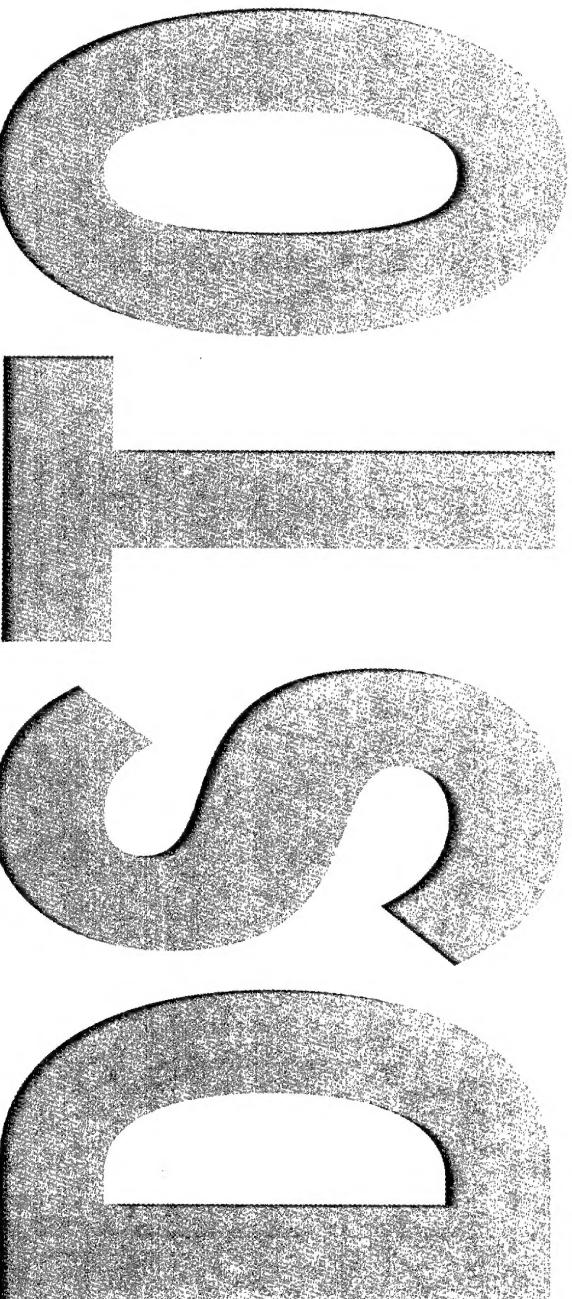




**Australian Government
Department of Defence**

Defence Science and
Technology Organisation

Aug 2003



Optimal Search, Location and Tracking of Surface Maritime Targets by a Constellation of Surveillance Satellites

Paul E. Berry, Carmine
Pontecorvo and David A.B. Fogg
DSTO-TR-1480

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

20040122 020



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Optimal Search, Location and Tracking of Surface Maritime Targets by a Constellation of Surveillance Satellites

Paul E. Berry, Carmine Pontecorvo, David A.B. Fogg

Intelligence, Surveillance and Reconnaissance Division
Information Sciences Laboratory

DSTO-TR-1480

ABSTRACT

The issues associated with the maximal exploitation of space-based surveillance resources are unique due to the nature of the platforms, their sensors and intermittent communications links to the ground station for target information processing and sensor tasking. This report applies the GAMBIT formalism [Berry & Fogg (1)] to the networked sensor decision problem of determining the optimal allocation of space-based surveillance resources for the purpose of detecting, locating and tracking surface maritime targets. As well as being an issue of interest in its own right, this application is an instance of the more general surveillance asset allocation problem.

RELEASE LIMITATION

Approved for public release

AQ F04-03-0220

Published by

*DSTO Information Sciences Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2002
AR-012-869
August 2003*

APPROVED FOR PUBLIC RELEASE

Optimal Search, Location and Tracking of Surface Maritime Targets by a Constellation of Surveillance Satellites

Executive Summary

The key to integrating surveillance operations and architectures is the allocation and tasking of multiple surveillance assets for the purpose of collectively satisfying specified surveillance information requirements subject to time, space and resource constraints. This is a classical networked sensor decision problem. A generic framework (GAMBIT) for addressing this issue was developed by Berry & Fogg [1], which assumes that the supporting technologies in terms of communications bandwidth and processing power are in place and are represented as information constraints. The issues associated with the maximal exploitation of space-based surveillance assets are unique due to the nature of the platforms, their sensors and the communications links to the ground station for the processing of target information and the tasking of the sensors. This report applies the GAMBIT formalism [1] to the problem of determining the optimal allocation of space-based surveillance resources for the purpose of detecting, locating and tracking surface maritime targets.

Authors

Paul Berry

Intelligence, Surveillance and Reconnaissance Division

Paul studied Mathematics at the University of London, majoring in mathematical physics, before pursuing a PhD on the topic of 'Waves on Flows with Vorticity' which he completed in 1979. For 12 years he was employed at research laboratories in the UK's electricity supply industry, working on fundamental scientific problems associated with the generation and transmission of electrical energy. During this period he became progressively more involved with the modelling and optimisation of large-scale engineering systems for the purposes of investment appraisal, design evaluation, economic operation and risk mitigation. He originally joined DSTO in 1991 to work on the stochastic analysis of telecommunication networks and moved to Surveillance Systems Division in 1998 in order to assume responsibility for the assessment of integrated surveillance operations and architectures.

Carmine Pontecorvo

Intelligence, Surveillance and Reconnaissance Div

Carmine Pontecorvo completed his undergraduate degree in Electrical and Electronic Engineering in 1994 and a Ph.D. in image processing in 1998, both at The University of Adelaide. He then went to Samsung SDS Co. in Seoul, South Korea for a year working on a surgical simulator tool. In 2000 he joined the Surveillance Systems Division of DSTO working on space-related issues, in particular imaging satellites for surveillance and maritime tracking.

David Fogg

Intelligence, Surveillance and Reconnaissance Div

David Fogg is a graduate of the University of Adelaide, currently a Principal Research Scientist, with over 30 years experience in the field of modelling and assessment of military systems including surveillance systems, combat systems and command and control. He has also been involved in the development of synthetic environments and in signal processing, mathematical analyses and operations research. He is currently the Head of Surveillance Systems Modelling and Assessment Group which is responsible in conjunction with sensor domain experts for providing a capability to model military surveillance systems and to assess their

performance and integrated capability. In addition, he is the Science and Technology Capability Coordinator for Land Intelligence Surveillance and Reconnaissance.

Contents

1. INTRODUCTION	1
2. GENERAL STATEMENT OF THE PROBLEM.....	1
2.1 Assumptions	2
3. GENERAL QUANTITATIVE FORMULATION OF THE PROBLEM.....	3
4. PHASES OF A SURVEILLANCE OPERATION	5
5. PHASE I: SEARCH.....	6
5.1 Bayesian update	7
5.2 Target motion prediction.....	8
5.3 Sensor Control Decisions	8
6. PHASE II: LOCATE.....	10
6.1 Bayesian update for location estimation of multiple targets.....	11
6.2 Target Motion Update	12
6.3 Sensor Control Decisions	12
6.4 Single target case.....	13
7. PHASE III: TRACK	17
7.1 Bayesian update for track estimate of a single target.....	18
7.2 Target Motion Prediction	18
7.3 Sensor Control Decisions	19
8. TRANSITIONING BETWEEN PHASES	19
9. IMPLEMENTATION	20
9.1 Target Motion Models	20
9.1.1 Dead Reckoning.....	20
9.1.2 Gauss-Markov	21
9.1.3 Random Course Change.....	22
10. PERFORMANCE AND COMPARISON	24
10.1 General Variables	24
10.2 Target Motion	24
10.3 Sensor Swaths.....	26
10.4 Alternative Swath Selection Techniques	27
10.4.1 Random Selection	27
10.4.2 Maximum Probability	28
10.4.3 Maximum Sum of Probability	28
10.5 Results	28
11. DEMONSTRATION OF BENEFITS FOR ANALYTICALLY TRACTABLE EXAMPLE	32

12. CONCLUSIONS	37
13. RECOMMENDATIONS	38
14. REFERENCES.....	39
15. MATHEMATICAL NOTATION.....	40
APPENDIX A: DERIVATION OF A GENERAL EXPRESSION FOR THE EXPECTED ENTROPY OVER A SET OF OBSERVATIONS.....	42
APPENDIX B: INTERPRETATION OF THE EXPRESSION FOR THE EXPECTED ENTROPY	43
APPENDIX C: IMPLEMENTATION PROCEDURES.....	44
C.1. Discretized Tracking Procedure	44
C.2. Gauss-Markov Motion Update Procedure.....	45
C.3. Optimal Swath Selection Procedure.....	46
C.4. Expected Entropy Change Procedure	47
C.5. Bayesian Update Procedure.....	48

1. Introduction

The advancing capabilities of sensors, continued miniaturisation of satellite payloads and reducing costs of satellite launch make surveillance from space an increasingly attractive option. However, the effective utilisation of the surveillance resource provided by space-based platforms requires that the data originating from their sensors be processed appropriately and that the sensors themselves, which provide the data, be tasked appropriately. For optimal resource utilisation, the retrospective exploitation of information and the prospective management of sensors need to be integrated into a single procedure.

This is a specific instance of the more general integrated surveillance problem, which is to determine how to allocate and task multiple surveillance assets to meet specified surveillance information requirements subject to time, space and resource constraints, while achieving efficiency and taking into account uncertainty arising from environmental effects, unpredictable target motions and imperfect sensors. Thus many of the practical implementation issues are addressed in this report and can be extended to the general problem.

An optimal formulation of the problem based upon Bayesian estimation and an entropy measure of surveillance effectiveness has been stated and a procedure proposed called GAMBIT (Berry & Fogg [1]). This was developed for the purpose of integrating a range of disparate surveillance sensor systems but is applied here to the specific problem of managing the sensors on a constellation of surveillance satellites for search operations (i.e., target detection), target location estimation and tracking of targets in the maritime domain. Bayesian approaches to tracking are gradually increasing in popularity (Stone *et al* [2] and Kastella [3]) and will continue to do so due to the emergence of efficient particle filtering techniques for Bayesian estimation. A number of recent papers have begun to explore the use of entropy for target tracking such as Xianxing *et al* [4], Fogg [6] and Xianxing *et al* [7].

2. General Statement of the Problem

The problem of tracking maritime surface targets from space is unique in the field of target tracking because of the fact that sensors do not scan an area at equally spaced intervals of time. Thus, sensor data can only be downloaded for processing during specific windows of opportunity when a satellite has access to a ground station. Furthermore the satellite sensors can only be tasked during such access periods. Hence there are delays in processing and tasking, which, though predictable, have to be taken into account in the tasking computations. This assumes, of course, a particular system architecture in which processing and tasking is performed on the ground. For architectures involving on-board processing and direct communication between surveillance satellites the delays could be reduced.

By *tasking* of a satellite is implied the choice of sensor from a sensor suite, its look-angle and its mode of operation and hence resolution. The fact that the area of interest is generally quite large (especially for search operations), revisit times potentially quite long, target motions inherently unpredictable and the environment obscuring or attenuating target signals means that account must be taken of the uncertainty associated with information. Hence the GAMBIT formalism [1] would seem to be appropriate for this problem. It would also be desirable to take account of any prior information regarding the existence and nature of a specific target, or information about targets in general, and our ability to observe them, as this becomes available, in order to assist in the execution of a space-based surveillance operation e.g. *a priori* target normalcy data (such as shipping lanes), information originating from an intelligence source about the time of departure of a particular target of interest from a port, meteorological information on cloud cover (in the case of EO sensors), and target speed and direction estimates extracted from imagery or track and constrained by the presence of islands and land masses.

How a space-based surveillance operation is executed also depends crucially on what information is sought e.g. whether the objective is to determine the existence of all targets in a specific area, to track all targets entering a given area or crossing a given line, or to track and, ultimately, classify and identify all targets conforming to a particular set of characteristics expressed in terms of location, speed, direction, origin and destination. These sets of characteristics may be expressed as hypotheses defined probabilistically.

It may also be useful to maintain mutually exclusive hypotheses regarding a target for the primary purpose of eliminating it from consideration, but hedging one's bets by maintaining the option of admitting it at a later point in time. There is of course a computational cost in maintaining alternative hypotheses but it provides the flexibility to change beliefs in a rapidly changing situation. By *belief* is meant the choice of one hypothesis versus its competitors for the purpose of making a decision, which has a physical outcome, such as tasking a sensor. If no physical action is required then all possible hypotheses can be maintained within available computational resources. When a choice has to be made between hypotheses then it must be done rationally i.e., its consequence must maximise the expected utility of the information collected in the context of the overall surveillance information requirement. Sometimes a compromise will be possible and a decision will be selected because it will assist in discriminating between competing hypotheses.

2.1 Assumptions

A constellation of surveillance satellites is used to search for, locate and track a ship. By use of the term *locate* is meant an improved estimate of a previously detected target's location and which is sufficient for it to be tracked. The following assumptions are made:

- The satellites' sensors have a look angle that enables them to sweep a specified swath (within their field of regard)
- The satellites' sensors have discrete modes of operation, each yielding a particular resolution and hence swath width, from which can be deduced target probabilities of detection and false alarm
- A *detection event* could correspond to an actual target or could be a false alarm. An estimate of target location and velocity is associated with a detection and incorporates observation errors. The Bayesian estimation approach can incorporate the effect of land through a likelihood function by associating the existence of ships only with areas of ocean for detection.
- A target motion model is used to predict future target locations (actual and false) for the purpose of tasking satellite sensors on future fly-overs. This takes account of the likely behaviour of the target of interest, the vicinity of ports to where it may be heading, and the vicinity of land and of straits between islands. This is effectively a filter that eliminates targets that are unlikely to be of interest on the basis of their dynamical characteristics.
- For the future fly-over of a satellite and predicted target locations, its sensor is tasked optimally to collect the required information subject to the sensor capability constraints (incorporating environmental effects)
- In order for observations to be utilised, account is taken of the need to communicate the data from a satellite to a ground station, to process it and to communicate it to the appropriate satellite from a ground station.

Three cases are proposed: a deterministic model of the target motion and sensor tasking, a Gauss-Markov target motion model and a conservative target motion model. Note that assumptions regarding the capabilities of sensors may be relaxed but in general will relate to the allowable size and shape of an area of coverage on a fly-pass and the corresponding resolution. The filtering of information in accordance with hypotheses for targets of interest can be done for multiple hypotheses concurrently. However, only one sensor action can be selected although this can be optimised for just one of the hypotheses or can aim to compromise by partially satisfying several hypotheses.

3. General quantitative formulation of the problem

The region \mathfrak{R} of interest is discretised into numbered cells, each cell of which has sufficiently small area that the probability of it being occupied by two or more independent targets is negligible. For simplicity the cells may be defined by the imposition of a rectangular grid upon \mathfrak{R} . Without loss of generality, let there be N cells where the probability of a cell being occupied by a target at a time t is $p_i(t)$, $i = 1, \dots, N$. Let the surveillance satellites have access to the region \mathfrak{R} at discrete instants of time $t_1, t_2, \dots, t_n, \dots$, known in advance and which will not, in

general, be equally spaced, then we write $p_i^n \equiv p_i(t_n)$. These times are referred to as epochs and let the area accessed by a satellite at the n^{th} epoch be A_n , where $A_n \subset \mathfrak{R}$. Note that p_i^n may be interpreted to mean either the probability of cell i containing a target at epoch n (as opposed to it not containing a target), or alternatively the probability that a single target in the region \mathfrak{R} occupies cell i (as opposed to any other cell). These represent alternative formulations of the problem and which of them is relevant will depend upon the information sought, as will be discussed.

The sensors' performances are modelled in terms of their probabilities of detection, p_d , and false alarm, p_{fa} . These probabilities will vary depending upon the resolution of the sensor mode used and the intervening environment. Note that the false alarm probability per cell has to be calculated from the false alarm distribution P_{fa} per unit area (i.e., a Poisson distribution in 2-dimensions) and the area of the cell i.e., $p_{fa} = P_{fa} \Delta A$ where ΔA is the area of a cell. For simplicity it is assumed that all cells are equal in area and that the false alarm distribution is uniform, hence p_{fa} is spatially constant for a given choice of resolution. The probability distribution is updated based upon the observations using Bayes' rule using a likelihood function appropriate to the way in which the state space is structured. A likelihood function, as defined in statistics, is the probability of obtaining a particular set of measurements for a given state of the system (assuming discrete states).

A target motion update is computed for predicted target motions between epochs n and $n+1$. Note that within this cell-based formulation, cells are assumed to be independent and no unique association is made between measurements and targets. This is adequate for determining the existence of targets within \mathfrak{R} in the search phase of a surveillance operation.

If one begins to associate observations uniquely with a particular target of interest either because it is known for a fact that there is only one target in \mathfrak{R} or because the means exists to distinguish that target from other targets in \mathfrak{R} , then one can specify the probability distribution for the location of a target given that one and only one such target exists in \mathfrak{R} . This can be done by conditioning the distribution upon the hypothesis that exactly one target exists. There will still be location errors and uncertainty about its location caused by false alarms but these effects will be captured by the probability distribution.

It is easier to condition the cell-based distribution upon the hypothesis of the existence of one target than perform Bayesian updates and target motion updates upon the conditioned probability distribution. It can be shown that results obtained are equivalent [1].

This probabilistic formulation facilitates a quantitative approach to area search for initial target detection, detected target location estimation and target tracking that is more general than traditional search and tracking techniques employing Kalman

filtering and assuming linear observation and target motion models with Gaussian errors and target accelerations. This is made possible by the fact that there will generally be adequate time between observations by satellite sensors to perform optimal detection and estimation using nonlinear filtering. For real-time tracking of targets with terrestrial radars this approach is not yet considered to be feasible although particle filtering techniques for nonlinear filtering are currently an area of active research.

This formulation also facilitates the optimal control of the satellite sensors in a predictive sense in order to maximise the expected information. This requires a quantitative measure of information for which we employ *information entropy* as discussed in Berry & Fogg [1].

4. Phases of a Surveillance Operation

It is assumed that a surveillance operation is broken down into separate search, locate and track phases, which are non-overlapping in time and have distinct transitions in time between them. The sensor management strategy is to be separately optimised within each phase so as to satisfy its distinct surveillance information requirements. It would be possible, in principle, to undertake such activities concurrently using the same set of surveillance assets with appropriate prioritisation of information, but this more general case is not considered here.

For each of the phases the information sought is different and hence the state space has to be structured appropriately. In the case of the search phase it is simply the existence of targets within \mathfrak{R} which is of interest. The locations of the targets are of interest only for the purpose of maximising search efficiency by minimising multiple detections of the same target. In the location estimation phase it is the locations of particular targets of interest that are of primary importance and hence the state space is structured to reflect this. In the tracking phase a single target of interest is tracked on the basis that observations made can be distinguished from those of other targets either on the basis of sufficient spatial separation or differing target characteristics.

For each of the phases of the surveillance operation the state spaces and probability distributions are defined to be consistent with the surveillance information requirement. Bayes' rule is applied using an appropriate likelihood function to update the probability distributions in the light of sensor-derived information about the existence and locations of targets. This is developed in the following sections.

For each phase of the surveillance operation, we select sensor actions with the objective of maximising the expected information, where *expectation* is used in the statistical sense. In general this could involve directing sensors and their platforms over a future infinite time horizon that, for practical purposes, would have to be made finite (see

Berry [5]). The problem is simplified in the case of a constellation of satellite sensors because we do not have control over the platforms (i.e., fixed constellations) and we have the option of delaying future tasking of sensors until the latest information is available. Hence we only need to plan for the next epoch following an observation at the previous epoch.

For a given phase the information requirement is specified in terms of the state space as described previously and information represented in terms of a probability distribution over the state space. The satisfaction of the information requirement, or equivalently the completeness of the information, is represented quantitatively by the entropy of the distribution. The objective in any phase of a surveillance operation is to reduce information uncertainty by minimising the entropy of the probability distribution representing it. Since one cannot predict precisely what the consequence of an observation will be following a sensor control action, the objective chosen is the *expectation* of the entropy following the future observations.

5. Phase I: Search

It is assumed that a surveillance operation begins with a search activity, whose objective may be to find all, some or just one target within a given area based upon some prior information as to the targets' locations or distribution. It could, in fact, have the objective of confirming that there are no targets in a given area. In general we will consider that the search phase of a surveillance operation has as its objective the determination of the existence or non-existence of targets within a region.

The most general formulation of the problem is obtained by defining a joint probability distribution for the numbers of targets in the region and their locations (Berry & Fogg [1]). This may be initialised using an *a priori* distribution, updated using Bayes' rule in response to observations at specific instants of time and evolved over time to reflect target motion. This is the most general formulation for the representation of targets and their locations in the case of targets not being distinguished, that is a measurement is not associated with any of a set of previous measurements for a target as is the case in target tracking. It is the most general formulation in the sense that it allows for correlations between targets (e.g. targets comprising a naval fleet) and the likelihood of any hypothesis expressed in terms of the existence, numbers and locations of targets in the region or a sub-region can be deduced from it. However, because the state space grows combinatorially with target numbers it is computationally inefficient to perform Bayesian updates and predict expected entropies arising from proposed sensor actions. This most general formulation of the search problem is referred to as the *global formulation*.

The reason why it is necessary to define a probability distribution over the region (as opposed to simply associating probabilities with detections) is because during a search

activity one is concerned with the *existence* of targets rather than their numbers or locations. It is therefore necessary to monitor those parts of a region that have been searched and when in relation to likely target speeds due to the possibility of targets entering an area previously searched.

A computationally more efficient formulation of the search problem [1] arises from the assumption that targets are independent of each other and their distribution conforms to a 2-D Poissonian spatial distribution that becomes a 2-D Bernoulli distribution when discretised. This is implemented by discretising the region into a grid, within each cell of which there is defined a probability that it contains a target. The cell size is taken to be sufficiently small that the probability of a cell containing more than one target is negligible. All cells are considered to be independent, and values for their probabilities are updated independently. Hence each cell is considered to be a state space in its own right with the probabilities of a target existing or not existing within it summing to unity. This is referred to as the *localised formulation* of the problem. There would be no attempt to track a target during the search phase of an operation since the objective is to associate a general area within the region with the existence of a target as quickly as possible. Location estimates will tend to degrade over time as a consequence due to target motion. Hypotheses can be derived directly from the localised formulation. It can be shown that rather than update probabilities conditioned upon hypotheses, updating probabilities in the localised formulation and then conditioning is equivalent and easier to implement [1].

The probabilities for cells within the swath A_n are subjected to Bayesian updates based upon whether or not a detection occurs in each of them. The probabilities corresponding to cells outside of A_n are not changed since no new information is received (i.e., no detections or non-detections). It can be demonstrated [1] that the entropy for the global distribution of targets is equal to the sum of the entropies for the individual cells in the localised formulation.

A sensor action is selected which minimises the expected entropy of the information sought, for this will, on average, most closely satisfy the information requirement. The expectation is performed on the future information based upon the current information. The sensor control action is assumed to be the sensor look angle α , the value of which determines uniquely the area of coverage $A(\alpha)$.

5.1 Bayesian update

If the *a priori* probability for a target in cell i at epoch $n+1$ is \hat{p}_i^{n+1} , and a detection occurs in that cell as a consequence of it being observed, then application of Bayes' rule gives the *a posteriori* probability

$$p_i^{n+1}(1) = \frac{p_d \hat{p}_i^{n+1}}{p_d \hat{p}_i^{n+1} + p_{fa} (1 - \hat{p}_i^{n+1})}$$

If a detection does not occur as a consequence of the cell being observed then

$$p_i^{n+1}(0) = \frac{(1-p_d)\hat{p}_i^{n+1}}{(1-p_d)\hat{p}_i^{n+1} + (1-p_{fa})(1-\hat{p}_i^{n+1})}$$

Cells that are not observed are not updated, as cells are assumed independent in the localised formulation.

5.2 Target motion prediction

This gives the *a priori* distribution \hat{p}_i^{n+1} at time t_{n+1} ($(n+1)^{\text{th}}$ epoch) given the *a posteriori* distribution p_i^n at time t_n due to non-deterministic target motion.

A Markovian target motion model is defined as follows. Assume that the probability that a target is in cell j at time t_{n+1} given that it was in cell i at time t_n is $Q(i, j, t_{n+1} - t_n)$.

Then

$$\sum_{j=1}^N Q(i, j, t_{n+1} - t_n) = 1$$

as the target that is in cell i at time t_n has be somewhere at time t_{n+1} (Conservation of Probability). It will eventually be necessary to take account of boundary conditions due to targets drifting in and out of the region \mathfrak{R} . Then the probability of a target being in cell j at time t_{n+1} due to possibly moving there from all other cells since time t_n is

$$\hat{p}_j(t_{n+1}) = \sum_{i=1}^N Q(i, j, t_{n+1} - t_n) p_i(t_n).$$

This target motion model is Markovian because the motion of the target depends purely upon where it was located at the last epoch and not on where it was located prior to that time. Because the times between epochs are unequal in length, the probability transition matrix Q is made explicitly dependent upon the time interval $t_{n+1} - t_n$.

Often a Gaussian distribution is chosen to represent the non-deterministic motion of a target between observations, in which case this would be referred to as a Gaussian Markov target motion model and Q would be defined appropriately. It will be seen when a target is being tracked its velocity as well as location need to be estimated and the state of a target at any time no longer depends only upon location. However, its target motion model may still be considered Markovian by extending its state space as will be seen. Target motion models are detailed in Section 9.1.

5.3 Sensor Control Decisions

It is necessary to determine which area $A_{n+1} \subset \mathfrak{R}$ a satellite should access at the $(n+1)^{\text{th}}$ epoch to maximise the information obtained. In order to specify this as a

quantitative optimisation problem we use information entropy as a measure of the quality of surveillance information. Entropy is defined by the probabilities over a state space. If we take the joint probability distribution for the numbers and locations of all possible targets in the localised formulation and take its entropy, then it can be shown [1] that the entropy is simply the sum of the entropies for the individual cells. Consequently one can consider the consequences of a detection or non-detection for each cell separately for each cell inspected. One can predict the contribution to the total entropy by the i^{th} cell immediately before any observations as

$$\hat{h}_i^{n+1} = -\hat{p}_i^{n+1} \log \hat{p}_i^{n+1} - (1 - \hat{p}_i^{n+1}) \log(1 - \hat{p}_i^{n+1})$$

using the *a priori* probabilities \hat{p}_i^{n+1} obtained from the target motion update. If one were to assume a perfect sensor then the consequence of inspecting the i^{th} cell, whether a detection occurred or not, would be to reduce its contribution to the total entropy to zero. In this idealised case the choice of area would be that for which the total contribution of all its cells before observation was a maximum, which would be obtained simply by choosing the swath with the most cells. This would require the assumption that $p_d = 1$ and $p_{fa} = 0$, that is a perfect sensor with no obscuration. For those cells not inspected there is no change to their contribution to the entropy due to the independence assumption.

If we are to accommodate the issue of sensor resolution through values of $p_d < 1$ and $p_{fa} > 0$ then we need to look at the *expected* change in entropy following an observation. Independence of cells enables this to be computed easily per cell.

Let $p_i^{n+1}(0)$ be the probability at the $(n+1)^{\text{th}}$ epoch after a hypothetical non-detection in the i^{th} cell and $p_i^{n+1}(1)$ be the probability after a hypothetical detection there. Then the entropy in each case is

$$h_i^{n+1}(0) = -p_i^{n+1}(0) \log p_i^{n+1}(0) - (1 - p_i^{n+1}(0)) \log(1 - p_i^{n+1}(0))$$

and

$$h_i^{n+1}(1) = -p_i^{n+1}(1) \log p_i^{n+1}(1) - (1 - p_i^{n+1}(1)) \log(1 - p_i^{n+1}(1))$$

and hence the *expected* entropy is

$$E(h_i^{n+1}) = \Pr\{0\}h_i^{n+1}(0) + \Pr\{1\}h_i^{n+1}(1)$$

where the probability of a detection is

$$\Pr\{1\} = \Pr\{\text{detection}\} = p_d \hat{p}_i^{n+1} + p_{fa} (1 - \hat{p}_i^{n+1})$$

and the probability of a non-detection is

$$\Pr\{0\} = \Pr\{\text{no detection}\} = 1 - \Pr\{\text{detection}\} = 1 - p_d \hat{p}_i^{n+1} - p_{fa} (1 - \hat{p}_i^{n+1})$$

The choice of control action α is that which maximises the expected change in entropy

$$\max_{\alpha} \sum_{i \in A(\alpha)} \{\hat{h}_i^{n+1} - E(h_i^{n+1})\}$$

where \hat{h}_i^{n+1} is the *a priori* entropy of the i^{th} cell at the $n+1^{\text{th}}$ epoch following target motion update.

This sum of cell entropies will never be reduced to zero in practice (corresponding to perfect information) because of targets moving between observations, inadequate coverage and imperfect sensors, so the Law of Diminishing Returns applies. At some point in time therefore additional surveillance effort will result in no net improvement in the quality of information. This point may never be reached if the objective of a surveillance operation is to search for a target (or targets) of a particular type and then locate them with sufficient accuracy to track them. Similarly if Phase I is limited in time. Assuming that Phase I is not unlimited in time then a condition will be tested to determine when Phase II begins.

Of particular interest may be the total number of targets in \mathfrak{R} , irrespective of their locations, in which case this parameter should be estimated. The expected, or average, number of targets in \mathfrak{R} at the n^{th} epoch is simply

$$\sum_{i=1}^N p_i^n$$

The distribution for the number of targets in \mathfrak{R} is

$$\begin{aligned} \Pr\{k \text{ targets in } \mathfrak{R}\} &= \sum_{j_1=1}^N \sum_{j_2=j_1+1}^N \dots \sum_{j_k=j_{k-1}+1}^N p_{j_1} p_{j_2} \dots p_{j_k} \prod_{\substack{j=1 \\ j \neq j_1 \\ \text{etc.}}}^N (1-p_j) \\ &= \sum_{j_1=1}^N p_{j_1} \sum_{j_2=j_1+1}^N p_{j_2} \dots \sum_{j_k=j_{k-1}+1}^N p_{j_k} \prod_{\substack{j=1 \\ j \neq j_1 \\ \text{etc.}}}^N (1-p_j) \end{aligned}$$

from which the entropy can be computed, if required.

6. Phase II: Locate

The search phase of a surveillance operation will continue until there is sufficient confidence that the number of targets in the region is known by testing the entropy for the hypothesis for target numbers. This will never reduce to zero due to false alarms, non-detections and target motion but at some point a decision has to be made to advance to the locate phase of a surveillance operation. This decision has to be made with reference to the overriding mission objective and the desire to minimise the entropy of the surveillance information requirement associated with it.

The objective of the locate phase is to determine the current locations of targets whose existence has previously been determined through execution of the search phase. These targets, once located, may be tracked in the next phase, enabling other sensors to be cued to identify or classify them as required. The probability distribution for the locate phase is the distribution for the target locations conditioned upon the hypothesis for their numbers in the region. When the entropy for this distribution is sufficiently small due to the availability of improved estimates for target locations, their locations may be replaced by separate spatial distributions for each target if they are distinguishable, or more simply by point estimates with Gaussian errors as is traditionally done in tracking.

We define the probability distribution for the locations of k targets in \mathfrak{R} at the n^{th} epoch thus:

$$P^n(j_1, j_2, \dots, j_k) = \frac{\Pr\{\text{targets in cells } j_1, j_2, \dots, j_k \mid k \text{ targets in } \mathfrak{R}\}}{\Pr\{k \text{ targets in } \mathfrak{R}\}}$$

where the numerator is obtained from the localised (cell-based) formulation using

$$\Pr\{\exists k \text{ targets in cells } j_1, \dots, j_k\} = p_{j_1}^n p_{j_2}^n \dots p_{j_k}^n \prod_{\substack{j=1 \\ j \neq j_1 \\ \text{etc}}}^N (1 - p_j^n)$$

and the denominator has previously been given in Section 5.3. The sum of the probabilities over all possible configurations of k targets is unity, i.e.,

$$\sum_{j_1 \dots j_k} P^n(j_1, \dots, j_k) = \sum_{j_1=1}^N \sum_{j_2=j_1+1}^N \dots \sum_{j_k=j_{k-1}+1}^N P^n(j_1, \dots, j_k) = 1.$$

Note that the distribution for the locations of k targets has ${}^N C_k$ states, which can become large for more than a few targets.

6.1 Bayesian update for location estimation of multiple targets

Rather than maintain a distribution for $P^n(j_1, j_2, \dots, j_k)$ and update it directly using Bayes' rule for the assumed number k of targets (which requires specification of the appropriate likelihood function), it is simpler to continue updating the individual cell probabilities separately (assuming cell independence as in the localised formulation) and then condition on the number of targets as shown in the previous paragraph. This is possible because the ordering of hypothesis conditioning and Bayesian updating is immaterial. The advantage of this is that because the localised formulation has a smaller number of states, fewer Bayesian updates have to be performed (N as opposed to ${}^N C_k$).

6.2 Target Motion Update

Again, it is easier to perform the target motion update in the localised formulation, as in Section 5.2, and then condition the probabilities on the hypothesis for the number of targets than to apply the target motion update directly to $P''(j_1, j_2, \dots, j_k)$.

6.3 Sensor Control Decisions

The entropy corresponding to the information regarding the locations of an assumed number k of targets at the n^{th} epoch is

$$h'' = - \sum_{j_1 \dots j_k} P''(j_1, j_2, \dots, j_k) \log P''(j_1, j_2, \dots, j_k)$$

Computation of this expression results in a measure of the quality of the information obtained about the k targets. The entropy for the hypothesis that there exist exactly k targets should also be monitored to ensure that the hypothesis is still relevant. In practice multiple hypotheses can be maintained and monitored, but only one physical sensor action can be selected. The choice of sensor action should be to prove or disprove a hypothesis (e.g., that there are k targets, as opposed to any other number) and to maximise the information associated with a hypothesis believed to be true (e.g., where the assumed k targets are located). There will almost always be some uncertainty associated with a hypothesis unless, or until, its entropy is exactly zero.

Entropy may provide an objective measure of information but in practice the large number of potential states makes it difficult to compute. However most states are likely to have very small probabilities associated with them and so need not be considered in the summation. The primary significance of this formulation is that it provides an objective, quantitative approach from which rational approximations may be derived for implementation.

The choice of sensor control parameter should be based upon the *expectation* of the entropy following the observation, as before. Whereas for each cell observed only two outcomes are possible for the case of independent cells (a detection or non-detection), for the present case in which the cells are not independent since they are constrained by the condition that there are a fixed number of targets, all possible configurations of detections and non-detections resulting from a choice of sensor control action (swath and resolution) have to be considered in performing the expectation of the entropy. Then the sensor control action that yields the best improvement in expected entropy has to be found and selected. If the number of targets were small and probabilities of false alarm small then it would certainly be practical to compute. This is demonstrated for the case of a single target ($k=1$) in the following section.

The optimal swath selection procedure can be found in Appendix C.3.

6.4 Single target case

In the special case of a single target, a likelihood function can be written for updating the probability distribution. We re-use the notation for the independent-cell case but interpret the probability distribution p_i^n to mean the probability that the single target occupies the i^{th} cell at the n^{th} epoch so that

$$\sum_{i=1}^N p_i^n = 1.$$

Suppose that at a particular epoch there are M cells under the swath A and a set of k detections occurring in cells $\{j_1, j_2, \dots, j_k\}$, one or none of which may correspond to the true target, the remainder being false alarms. Then, since we are assuming that all observations are associated with one and the same target, the probabilities over the cells comprising \mathfrak{R} have to be updated using Bayes' rule everywhere, whether a cell is under the swath or not. That is, because observing a target known to exist somewhere reduces the probability of it being elsewhere.

Bayes' rule updates the probability distribution for the target based on a set of observations in cells $\{j_1, j_2, \dots, j_k\}$ as follows:

$$p_i = \frac{L(j_1, j_2, \dots, j_k | i) \hat{p}_i}{\sum_{l \in \mathfrak{R}} L(j_1, j_2, \dots, j_k | l) \hat{p}_l}$$

where \hat{p}_i is the *a priori* distribution following the target motion update. In practice, of course, the probability of false alarm will be relatively low and the number of false alarms per sweep or scan will most likely be small in number.

The likelihood function used in Bayes' rule has different expressions depending upon whether cell i is in A , or not in A (L_{III}), and on whether, if it is in A , it is one of the $\{j_1, j_2, \dots, j_k\}$ (L_{II}) or not (L_{I}). Note that all of $\{j_1, j_2, \dots, j_k\}$ must be in A (because a sensor can only obtain detections where it observes). The Bayesian update procedure can be found in Appendix C.5.

Case I corresponds to the situation in which the target is contained in A but does not belong to the set of detections. So the likelihood function is the probability that the target is not detected but that there are k false alarms in A (and hence $M-k-1$ non-false alarms).

Case I: $i \in A$ but $i \notin \{j_1, j_2, \dots, j_k\}$ for $k = 0, \dots, M-1$

$$L_{\text{I}}(j_1, j_2, \dots, j_k | i \in A \cap i \notin \{j_1, j_2, \dots, j_k\}) = (1 - p_d) p_{fa}^k (1 - p_{fa})^{M-k-1}$$

Case II corresponds to the situation in which the target is contained in A but *does* belong to the set of detections. So the likelihood function is the probability that the

target *is* detected and that there are $k-1$ false alarms in A (and hence $M-k$ non-false alarms).

Case II: $i \in A$ and $i \in \{j_1, j_2, \dots, j_k\}$ for $k = 1, \dots, M$

$$L_H(j_1, j_2, \dots, j_k \mid i \in A \cap i \in \{j_1, j_2, \dots, j_k\}) = p_d p_{fa}^{k-1} (1 - p_{fa})^{M-k}$$

Case III corresponds to the situation in which the target is not in A , so all detections must be false alarms and the question of whether the actual target is detected or not does not arise. Hence the likelihood function is the probability of k false alarms and $M - k$ non-false alarms.

Case III: $i \notin A$ for $k = 0, \dots, M$

$$L_{III}(j_1, j_2, \dots, j_k \mid i \notin A) = p_{fa}^k (1 - p_{fa})^{M-k}$$

Note that all of the above assumes that the probabilities of detection p_d and false alarm p_{fa} (per cell) are spatially uniform in A and are known functions of the sensor resolution, which, together with the look angle, determine the size of A . The probabilities can in principle be made spatially and temporally non-uniform so to represent environmental effects.

The target motion update is performed as before to obtain the *a priori* distribution for the Bayesian update

$$\hat{p}_j(t_{n+1}) = \sum_{i=1}^N Q(i, j, t_{n+1} - t_n) p_i(t_n)$$

and the entropy evaluated over all states is

$$h^n = -\sum_{j=1}^N p_j^n \log p_j^n.$$

The following is an implementation of the expression for the expected entropy for the single target case with non-zero false alarm rate.

It is proven in Appendix A that the expectation of the entropy over all possible sets of observations $\underline{j} = \{j_1, j_2, \dots, j_k\} \in \Lambda$ at the $n+1^{\text{th}}$ epoch, where Λ is the set of all possible sets of observations, is

where $L(l|i)$ is the sensor's likelihood function defined as the probability of obtaining measurement set $\underline{l} = \{j_1, j_2, \dots, j_k\}$ given that the target is in cell i , \hat{p}^{n+1} is the *a priori* probability that the target is in cell i (following target motion update), and $P^{n+1}(l)$ is the probability of obtaining the measurement set \underline{l} . Hence

$$P^{n+1}(\underline{l}) = \sum_{i \in \mathfrak{N}} L(\underline{l} | i) \hat{p}_i^{n+1}.$$

The likelihood function depends upon the area A swept by the sensor and hence so does the expected entropy. It also depends upon the resolution of the sensor because this affects both the probability of detection and probability of false alarm. The intention should be to choose the area of coverage and sensor resolution so as to minimise the expected entropy. This will involve a trade-off between area of coverage and sensor resolution. There is complication caused by the fact that, although we are assuming a single target, there are potentially a large number of detections. However the false alarm rate will generally be small, which will tend to limit the number of detections occurring in practice. It is possible to interpret the meaning of the terms in equation 6.1 and this is done in Appendix B.

The difficulty in evaluating the expression for the expected entropy lies in the fact that there are summations over all possible numbers and configurations of observations due to false alarms, although most of these will have an extremely small probability. However, to totally ignore the effect of false alarms would be to ignore the ambiguities that could occur when the sensor resolution is low. That is, not only would the probability of detection be reduced but also the possibility of obtaining false alarms could increase.

To make it clear that the expected entropy is the expectation over the numbers of detections and their configuration, we rewrite equation 6.1 by replacing \underline{l} by $\{j_1, j_2, \dots, j_k\}$. Then

$$\begin{aligned} &= h - \sum_{i \in \mathfrak{N}} p_i \sum_{k=0}^M \sum_{\{j_1 \dots j_k\} \in \Lambda} L(j_1 \dots j_k | i) \log L(j_1 \dots j_k | i) + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\} \in \Lambda} P(j_1 \dots j_k) \log P(j_1 \dots j_k) \\ &= h - \sum_{i \in A} p_i \sum_{k=0}^{M-1} \sum_{i \notin \{j_1 \dots j_k\}} L_I \log L_I - \sum_{i \in A} p_i \sum_{k=1}^M \sum_{\{j_1 \dots j_k\} \in i} L_{II} \log L_{II} - \sum_{i \notin A} p_i \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} L_{III} \log L_{III} \\ &\quad + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} P(j_1 \dots j_k) \log P(j_1 \dots j_k) \\ &= h - \sum_{i \in A} p_i \sum_{k=0}^{M-1} \sum_{i \notin \{j_1 \dots j_k\}} (1 - p_d) p_{fa}^k (1 - p_{fa})^{M-k-1} \log \{(1 - p_d) p_{fa}^k (1 - p_{fa})^{M-k-1}\} \\ &\quad - \sum_{i \in A} p_i \sum_{k=1}^M \sum_{i \in \{j_1 \dots j_k\}} p_d p_{fa}^{k-1} (1 - p_{fa})^{M-k} \log \{p_d p_{fa}^{k-1} (1 - p_{fa})^{M-k}\} \\ &\quad - \sum_{i \in A} p_i \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} p_{fa}^k (1 - p_{fa})^{M-k} \log \{p_{fa}^k (1 - p_{fa})^{M-k}\} \\ &\quad + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} P(j_1 \dots j_k) \log P(j_1 \dots j_k) \end{aligned}$$

where

$$\begin{aligned}
P(j_1 \dots j_k) &= \sum_{i \in \mathfrak{N}} L(j_1 \dots j_k | i) p_i \\
&= \sum_{\substack{i \in A \\ i \notin \{j_1 \dots j_k\}}} L_I p_i + \sum_{\substack{i \in A \\ i \in \{j_1 \dots j_k\}}} L_{II} p_i + \sum_{i \notin A} L_{III} p_i
\end{aligned}$$

has separate expressions for the cases $k=0$, $0 < k < M$ and $k=M$. For $k=0$:

$$P(\text{no detections}) = \sum_{i \in A} L_I p_i + \sum_{i \notin A} L_{III} p_i = (1-p_d)(1-p_{fa})^{M-1} p_A + (1-p_{fa})^M (1-p_A)$$

where $p_A = \sum_{i \in A} p_i$. For $0 < k < M$:

$$P(j_1, \dots, j_k) = p_{fa}^{k-1} (1-p_{fa})^{M-k-1} [(1-p_d)p_{fa}(p_A - \sum_{l=1}^k p_{j_l}) + p_d(1-p_{fa}) \sum_{l=1}^k p_{j_l} + (1-p_A)p_{fa}(1-p_{fa})]$$

and for $k=M$:

$$P(j_1, \dots, j_M) = \sum_{i \in A} L_{II} p_i + \sum_{i \notin A} L_{III} p_i = p_d p_{fa}^{M-1} p_A + p_{fa}^M (1-p_A)$$

The expected entropy expression may further be simplified thus:

$$\begin{aligned}
h' &= h - \sum_{i \in A} p_i \sum_{k=0}^{M-1} {}^{M-1}C_k (1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1} \log \{(1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1}\} \\
&\quad - \sum_{i \in A} p_i \sum_{k=1}^M {}^{M-1}C_{k-1} p_d p_{fa}^{k-1} (1-p_{fa})^{M-k} \log \{p_d p_{fa}^{k-1} (1-p_{fa})^{M-k}\} \\
&\quad - \sum_{i \in A} p_i \sum_{k=0}^M {}^M C_k p_{fa}^k (1-p_{fa})^{M-k} \log \{p_{fa}^k (1-p_{fa})^{M-k}\} \\
&\quad + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} P(j_1 \dots j_k) \log P(j_1 \dots j_k) \\
&= h - p_A \sum_{k=0}^{M-1} {}^{M-1}C_k (1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1} \log \{(1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1}\} \\
&\quad - p_A \sum_{k=1}^M {}^{M-1}C_{k-1} p_d p_{fa}^{k-1} (1-p_{fa})^{M-k} \log \{p_d p_{fa}^{k-1} (1-p_{fa})^{M-k}\} \\
&\quad - (1-p_A) \sum_{k=0}^M {}^M C_k p_{fa}^k (1-p_{fa})^{M-k} \log \{p_{fa}^k (1-p_{fa})^{M-k}\} \\
&\quad + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} P(j_1 \dots j_k) \log P(j_1 \dots j_k)
\end{aligned}$$

or,

$$\begin{aligned}
h' = & h - c_I(p_d, p_{fa}, M)p_A - c_{II}(p_d, p_{fa}, M)p_A - c_{III}(p_d, p_{fa}, M)(1-p_A) \\
& + \sum_{k=0}^M \sum_{\{j_1 \dots j_k\}} P(j_1 \dots j_k) \log P(j_1 \dots j_k)
\end{aligned}$$

It is apparent that the expression for the expected entropy is a series in ascending powers of p_{fa} and that we can decide to truncate the series at any term so as to simplify computation. For instance we could decide to retain terms in p_{fa}^0 and p_{fa}^1 but to discard terms in p_{fa}^2 and above on the grounds that they are negligibly small. Note that care is required here because although the probability of a particular pair of false alarms may be small, the probability of *any two* false alarms occurring somewhere may not be insignificant. In other words we have to take account of not just the magnitude of p_{fa}^k , but of ${}^M C_k p_{fa}^k$.

The procedure for computing the expected entropy change is in Appendix C.4.

7. Phase III: Track

The localised formulation of Section 5 has been used for tracking [3] but is, the authors believe, fundamentally flawed as well as grossly inefficient. This is because detecting a target with certainty at a location does not cause the probability of a target existing in its neighbourhood to collapse to zero, even if it was known that there was only a single target in that neighbourhood. This is a consequence of the fact that information is stored from a spatially-oriented perspective and not from a target-oriented perspective. In particular, if a measurement is known to be associated with a particular target but that fact is not recorded, or not capable of being recorded within one's information model, then information has effectively been discarded.

The way around this is to associate a measurement with a *particular* target and maintain separate probability distributions for the locations of each target. This process of associating measurements with targets is inherent to most trackers such as the PDA tracker [4] and there is little point in attempting to reproduce such a tracker here except to say that they are probabilistically-based and therefore consistent with the present GAMBIT formalism. Suffice it to say that we will assume for present purposes that the association between targets and measurements has been perfectly achieved. Note that if there is only a single target within the region of interest, or that there are multiple targets but measurements can be reliably associated with particular targets, then the issue does not arise.

7.1 Bayesian update for track estimate of a single target

Location estimation using Bayes' rule does not generate point estimates from which tracks can be composed; instead it produces a probability distribution for the location of the target over the region \mathfrak{N} . To extract point estimates from the probability distribution for track formation purposes would be to discard information and therefore lose the benefits of a Bayesian approach. We choose, therefore, to extend the Bayesian location estimation formulation to track formation so as to continue working with probability distributions and avoid having to extract point estimates. When probabilities for track states are found to be negligibly small then they can be discarded.

For present purposes a track is defined as a pair of locations representing both where the target is at the latest epoch and where it was at the previous epoch. We define a probability distribution over the state space of all such location pairs as follows

$$p_{i,j}^n = \Pr \{ \text{target moved from cell } i \text{ at epoch } n-1 \text{ to cell } j \text{ at epoch } n \}.$$

Applying Bayes' rule to obtain the probability at epoch $n+1$ from the probability at epoch n based upon a set of observations $\underline{l}^{n+1} = \{j_1, j_2, \dots, j_k\}$

$$\begin{aligned} p_{i,j}^{n+1}(\underline{l}^{n+1}) &= \frac{\Pr \{ \text{target moves from } i \text{ to } j \cap \text{observations } \underline{l}^{n+1} \}}{\Pr \{ \underline{l}^{n+1} \}} \\ &= \frac{L(\underline{l}^{n+1} | i, j) \hat{p}_{i,j}^{n+1}}{\Pr \{ \underline{l}^{n+1} \}} \end{aligned}$$

The likelihood function for the set of observations \underline{l}^{n+1} at epoch $n+1$ does not, in fact, depend upon the target location at epoch n , only on where the target is at the instant of observation at epoch $n+1$. Hence we can write

$$p_{i,j}^{n+1}(\underline{l}^{n+1}) = \frac{L(\underline{l}^{n+1} | j) \hat{p}_{i,j}^{n+1}}{\Pr \{ \underline{l}^{n+1} \}} = \frac{L(\underline{l}^{n+1} | j) \hat{p}_{i,j}^{n+1}}{\sum_{i \in \mathfrak{N}} \sum_{j \in \mathfrak{N}} L(\underline{l}^{n+1} | j) \hat{p}_{i,j}^{n+1}}$$

where $\hat{p}_{i,j}^{n+1}$ is the *a priori* probability following the target motion update, and the likelihood function $L(\underline{l}^{n+1} | j)$ is the same as the one previously defined for estimating target location in the case of a single target. The Bayesian procedure is presented in Appendix C.5.

7.2 Target Motion Prediction

A Markovian target motion model that exploits knowledge about estimates of the target velocity can be used to obtain the *a priori* distribution at epoch $n+1$ from the *a posteriori* distribution at epoch n . This is simply

$$\hat{p}_{i,j}^{n+1} = \sum_{i'} Q(j|i, i') p_{i',j}^n$$

where the transition probability matrix Q now depends not only upon where the target was at the previous epoch (cell i at time t_n) but how it got there (from cell i' at time t_{n-1}), and hence implicitly upon its velocity. The target motion model is therefore more general than that used for the search and location estimation phases since it takes into account not just the location of targets but estimates of their speed and direction. Note that Q is also implicitly dependent upon the time intervals $t_{n+1} - t_n$ and $t_n - t_{n-1}$.

To obtain the predicted distribution for location only we sum over all possible start locations i :

$$\hat{p}_j^{n+1} = \sum_i \hat{p}_{i,j}^{n+1} = \sum_{i',i} Q(j|i, i') p_{i',j}^n.$$

7.3 Sensor Control Decisions

The expected entropy depends purely upon the sensor, environment and target location but not on the target velocity. Hence the same expression may be used as for the single target location estimation phase except that better estimates of its location are used by virtue of the fact that the target motion model is implicitly adapted to the target's previous behaviour. Therefore one would expect the tracking algorithm to perform better than the single target location estimation algorithm.

8. Transitioning Between Phases

Since the optimal surveillance strategies are different for each of the phases and we assume that the phases are non-overlapping, a means for determining when to transition between them is required. A phase transition will occur when that phase has been "completed", i.e., when the required information for that phase has been obtained. So, for example, phase I ends when all targets have been detected and phase II ends when all detected targets have been located sufficiently accurately to be tracked. In reality, however, the surveillance information requirement will never be perfectly satisfied because the targets are moving and the information is continually degrading. A point will be reached at which additional effort makes little or no difference to the quality of information obtained. By quantifying the information obtained during each phase and comparing it against the information sought we have the means of monitoring the effectiveness of the surveillance and deciding when additional effort will result in diminishing returns. When this occurs, effort is better expended on the next phase of the surveillance operation.

The surveillance information requirement is specified as a probability distribution over a suitably defined state space and its effectiveness measure as the entropy of that distribution. The strategy selected for a given phase should aim to minimise the

information uncertainty and hence the entropy. When the entropy, or alternatively the rate of improvement of entropy, has reached a specified threshold then the decision may be taken to proceed to the next phase of the operation. When transitioning between phases one assumes that the preceding information requirement has been satisfied. So for example when transitioning from phase I to phase II it is assumed that all targets have been detected, and when transitioning from phase II to phase III it is assumed that all targets to be tracked have been located. Furthermore, we assume that no new targets enter the region of interest once phase I begins.

This is, in fact, the concept of a Bayesian network in which one reduces the complexity of a joint state space over all possible targets and their attributes to a set of interdependent hypotheses in order to manage the combinatorially large state space. This improves manageability but at the expense of reduced optimality. In practice there will be uncertainty associated with these assumptions so they may be treated as hypotheses and their probabilities monitored. So for instance in phase II, while the sensors are being utilised to accurately estimate the location of targets known to exist from phase I, they will not be revisiting areas in which targets were not found. Targets could infiltrate those areas over time so the target motion model will tend to reduce confidence in the hypothesis regarding the total number of targets in the region.

9. Implementation

Phase III of the surveillance operation was simulated in order to assess the potential improvements in tracking capability from implementation of the GAMBIT algorithm. The problem to which this technique is applied consists of a region of interest (ROI) with a single target moving diagonally across it; the ROI periodically falls within the field of regard of four non-overlapping swaths that notionally belong to one or possibly more satellites.

Phase I (search phase) has been addressed within a different operational context in Berry [5]. The procedure for implementing tracking is presented in Appendix C.1.

9.1 Target Motion Models

9.1.1 Dead Reckoning

If the target speed v is measured to within $\pm \Delta v$ and its heading θ to within $\pm \Delta\theta$, we assume it continues along a straight line path and estimate its position after a time t_i to be contained within a segment of an annulus between the circles

$$(x - x_i)^2 + (y - y_i)^2 = [(v - \Delta v)t_i]^2, \text{ and}$$

$$(x - x_i)^2 + (y - y_i)^2 = [(v + \Delta v)t_i]^2$$

and bounded by the two lines

$$\begin{aligned} (y - y_i) &= \tan(-\Delta\theta)(x - x_i), \\ (y - y_i) &= \tan(\Delta\theta)(x - x_i). \end{aligned}$$

9.1.2 Gauss-Markov

An explanation of the Gauss-Markov motion model in greater detail can be found in Fogg [6]. Let n denote the n^{th} epoch, and m_n be the number of time steps (Δt) between the n^{th} epoch and the $(n-1)^{\text{th}}$ epoch. Assume that the target is at position (x_0, y_0) at time zero. Furthermore, assume, that the x and y components of the position are independent, and the accelerations in each are white Gaussian sequences, ε_n , with the same variance ξ . Accelerations are constant over a time step Δt . Denoting the sequences of the x - and y -components of the target's position as $\{x_n\}$ and $\{y_n\}$, then these are Gauss-Markov and the equations of motion are 2nd order. For $\{x_n\}$ these equations are as follows (those for $\{y_n\}$ are similar):

$$\begin{aligned} x_{n+1} &= x_n + m_{n+1} \Delta t \dot{x}_n + \frac{1}{2} (m_{n+1} \Delta t)^2 \varepsilon_n \\ \dot{x}_{n+1} &= \dot{x}_n + m_{n+1} \Delta t \varepsilon_n \end{aligned}$$

where \dot{x}_n is the target's x -component of velocity at epoch n .

The conditional probability of finding the target at (x_n, y_n) at a time $m_n \Delta t$ after it was at (x_{n-1}, y_{n-1}) is denoted by $p\{(x_n, y_n) | (x_{n-1}, y_{n-1})\}$ and is known for this type of model to be bivariate Gaussian with x_n having mean

$$\mu_x = x_{n-1} + \dot{x}_{n-1} m_n \Delta t$$

and variance

$$\sigma^2 = \frac{4m_n^3 - m_n}{12} \xi \Delta t^4.$$

y_n has a similar expression for the mean and an identical variance [10]. Thus,

$$p\{(x_n, y_n) | (x_{n-1}, y_{n-1})\} = \frac{1}{2\pi\sigma^2} \exp\left[\frac{-1}{2\sigma^2} \left\{(x_n - \mu_x)^2 + (y_n - \mu_y)^2\right\}\right].$$

The value of ξ is a priori information, which is determined by the target's known motion over time relative to the chosen time increment Δt .

The Gauss-Markov Motion update procedure can be found in Appendix C.2.

9.1.3 Random Course Change

In this model a ship is assumed to continue on its estimated course at constant estimated speed with at most one change in course at a random chosen instant of time, with the course change angle chosen from a specified distribution. Note that this may seem to suggest that we begin with a deterministic start point and velocity (track) but the model is implemented for all possible start locations and velocities, weighted according to the probabilities calculated from the Bayesian update. We assume that the target is at the origin at observation time (epoch) t_n and that the next observation time is $t_{n+1} = t_n + T$. Without loss of generality we align the direction of motion with the Cartesian coordinate system as indicated in Figure 1 and assume the target speed is s .

Each point is uniquely specified by the pair of variables (x, y) or (z, θ) so there is a 1-to-1 mapping between them. In particular for any given point in the (x, y) -plane, there is only one way of getting to it from the origin in terms of distance z before a course change, and angle θ . Depending upon the distributions for z and θ , some points may be inaccessible. In particular points distant greater than sT from O are inaccessible in time T .

We begin by writing the transformation between (x, y) and (z, θ) . This is simply

$$\begin{aligned} x &= (sT - z) \sin \theta \\ y &= z + (sT - z) \cos \theta \end{aligned}$$

Next we derive the p.d.f. (probability density function) for the probability of reaching a point defined in terms of (z, θ) from O . We treat z and θ as random variables and suppose that they are independent. Then the joint p.d.f. for z and θ can be written

$$\Pr\{\text{reaching element } (z, z+dz) \times (\theta, \theta+d\theta)\} = p(z, \theta) dz d\theta = P_z(z) P_\theta(\theta) dz d\theta.$$

The function $P_\theta(\theta)$ may be arbitrary but should be chosen so as to reflect reality as closely as possible. It could even be spatially dependent to take account of straits between islands for example. The function $P_z(z)$ needs to be chosen with some care, however, because the time of an observation t_n is statistically uncorrelated with the time $t_n + z/s$ of a course change. The simplest assumption is to make the time between course changes negatively exponentially distributed. Thus

$$P_z(z) = \lambda e^{-\lambda z}$$

where λ may be interpreted as the average rate of course changes per unit distance.

Finally, the p.d.f. has to be converted into (x, y) coordinates. This requires the determinant of the Jacobian of the coordinate transformation as follows:

$$\begin{aligned}
 \Pr\{\text{reaching element } (x, x+dx) \times (y, y+dy)\} &= p(x, y) dx dy \\
 &= p(x(z, \theta), y(z, \theta)) \left| \frac{\partial(x, y)}{\partial(z, \theta)} \right| dz d\theta \\
 &= p(z, \theta) dz d\theta
 \end{aligned}$$

Hence

$$p(x, y) = \frac{p(z, \theta)}{\left| \frac{\partial(x, y)}{\partial(z, \theta)} \right|} = \frac{P_z(z) P_\theta(\theta)}{\left| \frac{\partial(x, y)}{\partial(z, \theta)} \right|}$$

where

$$\left| \frac{\partial(x, y)}{\partial(z, \theta)} \right| = \begin{vmatrix} \partial x / \partial z & \partial x / \partial \theta \\ \partial y / \partial z & \partial y / \partial \theta \end{vmatrix} = (sT - z)(1 - \cos \theta)$$

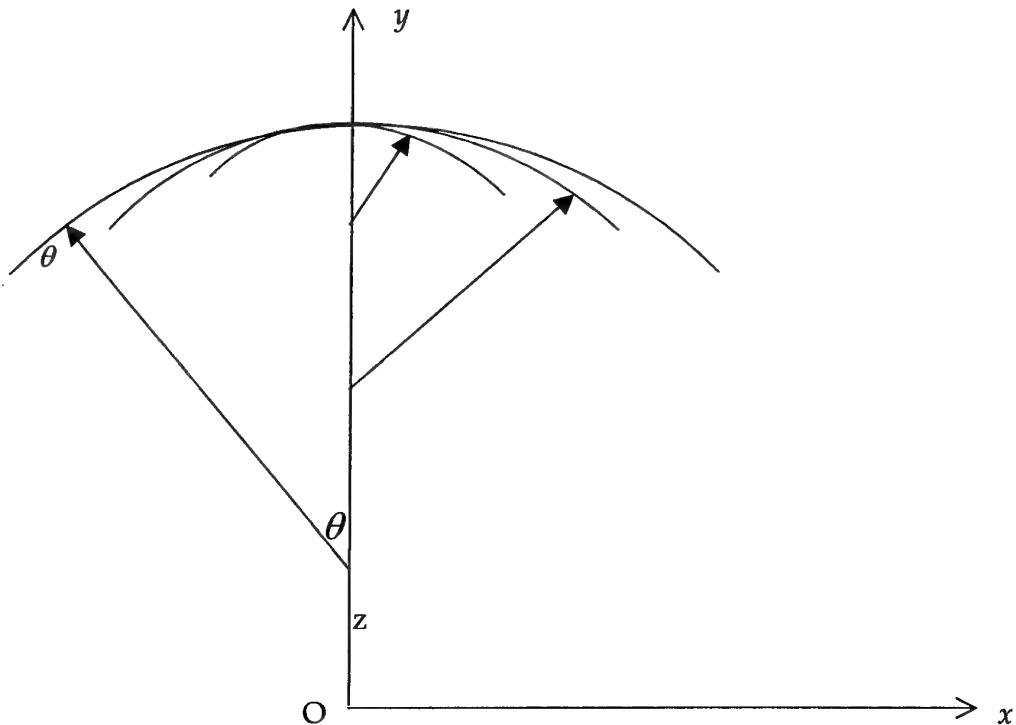


Figure 1: Possible ship trajectories

10. Performance and Comparison

10.1 General Variables

At this stage of the development of the Bayesian tracker, the algorithm is tested over a relatively small region of interest (ROI) due to its memory and computational requirements. If the ROI has $N \times N$ cells, then the total number of cells that must be stored is $N \times N \times N^2$ as where the target is expected to be for the next observation is conditioned on where it was after the last observation.

Ideally we want the ROI to span the entire area where a target could exist and also have a high resolution (i.e., large number of cells per given area target size), but a ROI satisfying both of these requirements would be impossible to store on computer. So we choose a ROI size of 16×16 enabling us to obtain Monte Carlo statistics on the performance of various tracking methods in a reasonable amount of time.

To reduce the effect of "motion-leakage" or "probability-leakage" across the boundaries of the ROI, two layers of "buffer" cells are placed around the entire perimeter of the ROI, allowing these probabilities to be accounted for, to a certain extent, beyond the area covered by the ROI. The probabilities within these buffer cells allow us to compute the probability flow of a target back into the ROI. For any "leakage" beyond these buffer cells, however, a single variable is used to retain the total probability that the target is neither in the ROI nor the buffer cells. This enables the Bayesian update to be computed correctly since all probabilities are accounted for. Thus, the overall selection of the size of the ROI is a compromise on having enough cells of sufficient resolution to cover a desired target area and also having enough buffer cells to minimize the effect of probability overflow, whilst minimizing the memory and computational demands made on the computer by having too many cells.

10.2 Target Motion

In order to test the basic utility of the expected entropy algorithm, a relatively simple target motion trajectory has been chosen, and is shown in Figure 2 below. The target begins at an (x, y) position of $(1,1)$, moves diagonally down until halfway, then makes a course change, and eventually moves diagonally down to the bottom left-hand corner cell. A second motion trajectory is shown in Figure 3, where the target continues beyond the halfway point, but now changes direction at cell $(10,10)$. This trajectory ensures that the target does not change direction at the boundary of two swaths, which are defined later.

We assume that the sensor swaths access the ROI once per time unit or epoch; the target speed is then chosen such that the target moves to a different cell for each epoch.

The motion is always linear except for one course change. This motion not only tests the ability of the various swath selection techniques to track the target as it moves in a linear fashion, but also whether they can retain or reacquire access to the target once it makes a drastic change in heading.

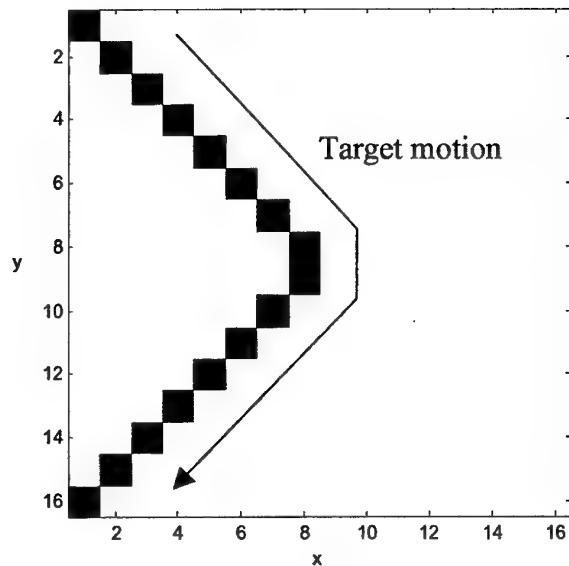


Figure 2: Target motion trajectory. The target moves from the upper left-hand corner to the lower left-hand corner.

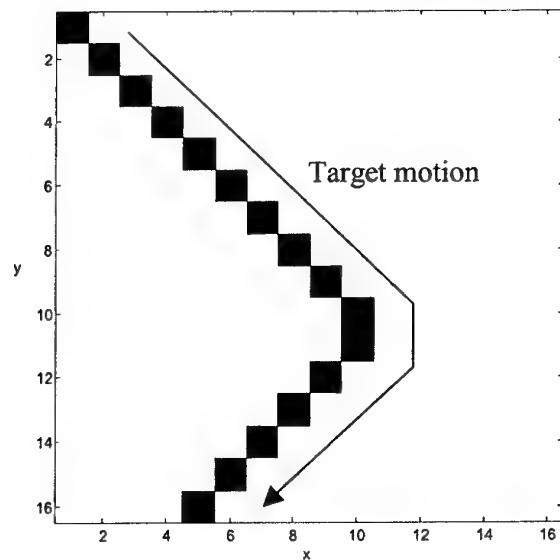


Figure 3: Second target motion trajectory. The target moves from the upper left-hand corner towards the lower left-hand corner.

10.3 Sensor Swaths

Throughout these experiments we use only 4 identical swaths as illustrated in Figure 4. These are non-overlapping but together cover the entire ROI. Each swath has a width of 4 cells and a length of 16 cells. We assume that the sensors are also identical, namely the probabilities of detection and false alarm.

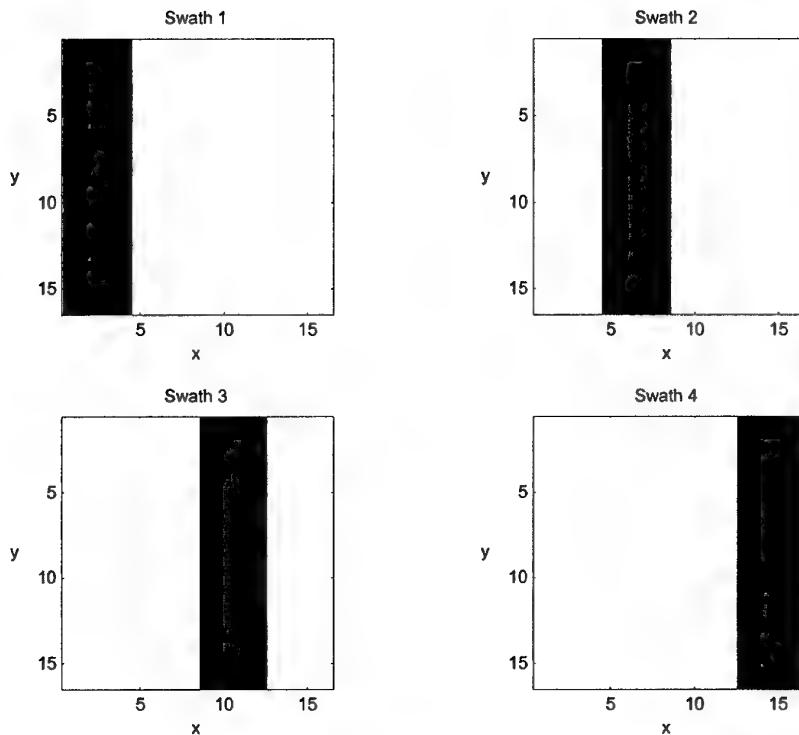


Figure 4: The four swaths (displayed as black strips) used in the experiment. Each is 4 cells wide and 16 cells long, with identical probabilities of detection and false alarm.

10.4 Alternative Swath Selection Techniques

This paper describes and derives a method of optimally selecting a swath for target tracking by choosing the swath for each access that minimizes the expected entropy across the ROI. Here, we discuss three other techniques that can be used to select the swath. They are the *random*, *maximum probability*, and *maximum sum of probability* methods.

10.4.1 Random Selection

This is the simplest swath selection technique. At each epoch, we randomly select one of the four swaths displayed in Figure 4. Thus, each swath has a $1/4$ chance of being selected during any access period or epoch.

The results from the random swath selection technique can be considered as the *baseline performance*, as it is the simplest possible selection technique if we had no information on where the target is and where it is heading. If we did have this

information then it is a simple matter to choose the appropriate swath, thereby resulting in excellent performance.

10.4.2 Maximum Probability

In this technique the swath that covers the cell with the largest probability is chosen. This is intuitive, and mathematically can be expressed as:

Select swath S_i such that: $\max_{j \in S_i} \{\bar{p}_j^{n+1}\}$ is largest for $i = 1, 2, 3, 4$.

10.4.3 Maximum Sum of Probability

This technique is similar to the maximum probability one, except that we choose the swath with the greatest sum of probabilities from all the cells lying under the swath. Thus,

Select swath S_i such that: $\sum_{j \in S_i} \bar{p}_j^{n+1}$ is largest for $i = 1, 2, 3, 4$.

The problem with the three techniques described above is that they do not take into account the non-idealities of the sensor i.e., they are independent of the sensor's probabilities of detection and false alarm. The Bayesian technique, however, can incorporate such sensor information.

10.5 Results

To run the simulation requires a number of variables to be set. These are presented in Table 1. The variable *Runs* refers to the number of Monte Carlo simulations performed to obtain the performance statistics. *N* is the width (and breadth) of the ROI, and *Buffer* is the number of layers of cells added to the perimeter of the ROI to reduce the effects of probability overflow. v_{x0} and v_{y0} are the target's initial speed values in the *x* and *y* directions. We do not assume that these values are known so they are set to 0. v_{\max} is the greatest possible speed that the target could achieve, and is used mostly to reduce the computational complexity of the problem. Δt is a fixed time interval over which the target's acceleration is constant, and *n* is the number of such time interval between successive accesses or epochs. Throughout all these experiments, $n\Delta t$ was kept at a constant value. The value of ξ determines the width of the Gaussian in the target motion model and is varied in the examples provided, as are the sensor characteristics of p_d and p_{fa} .

Table 1: Variables and the corresponding values used in the simulations.

Variable	Value
Runs	100

N	16
Buffer	2
v_{x0}, v_{y0}	0
v_{\max}	4
Δt	1/60
n	60
ξ	See Figures
p_d, p_{fa}	See Figures

To measure the performance of the various techniques we compute the total normalized entropy across all cells of the ROI. i.e.,

$$H = \frac{-1}{K} \sum_{j \in ROI} p_j \log p_j$$

where K is a normalizing factor corresponding to the maximum possible entropy, which occurs when the probability distribution is uniformly distributed across all cells. For $N \times N$ cells:

$$K = -\sum_{j=1}^{N^2} p_j \log p_j = -\sum_{j=1}^{N^2} \frac{1}{N^2} \log \left(\frac{1}{N^2} \right) = -N^2 \frac{1}{N^2} \log \left(\frac{1}{N^2} \right) = \log N^2.$$

Better performance is indicated by a smaller normalized entropy value over the ROI. From here on all entropy values quoted are the normalized values.

Figure 5 shows a comparison of the four possible swath selection techniques for the target motion as illustrated in Figure 2 for values of $\xi = 12.5$, $p_d = 0.95$ and $p_{fa} = 0.001$. The other parameters used in the simulation are given in Table 1. The entropy values are averaged over the total number of runs ("Runs").

Overall the random selection technique i.e., the baseline method, appears to be the worst, while the other three selection methods appear to give essentially the same, but much better, performance.

For the first few epochs the expected entropy and maximum sum of probability methods are worse than the other two, as they are affected by the probability of leakage out from the ROI and into the buffer cells. This leakage does not allow them to choose swath 1 (which would cover the target for the first few epochs). The random and maximum probability methods (for these epochs) are both unaffected by leakage, hence they can select swath 1, resulting in a detection event and smaller overall entropy.

Notice that there are two "spikes" at epochs 5 and 9, and to a lesser extent at epoch 13, due to the target probability distribution crossing over from one swath to another, hence the spreading of this distribution over two adjacent swaths. Selecting either swath does not cover the entire probability distribution; thus the entropy momentarily goes up, but all three eventually lock onto the target and reduce the entropy to almost zero.

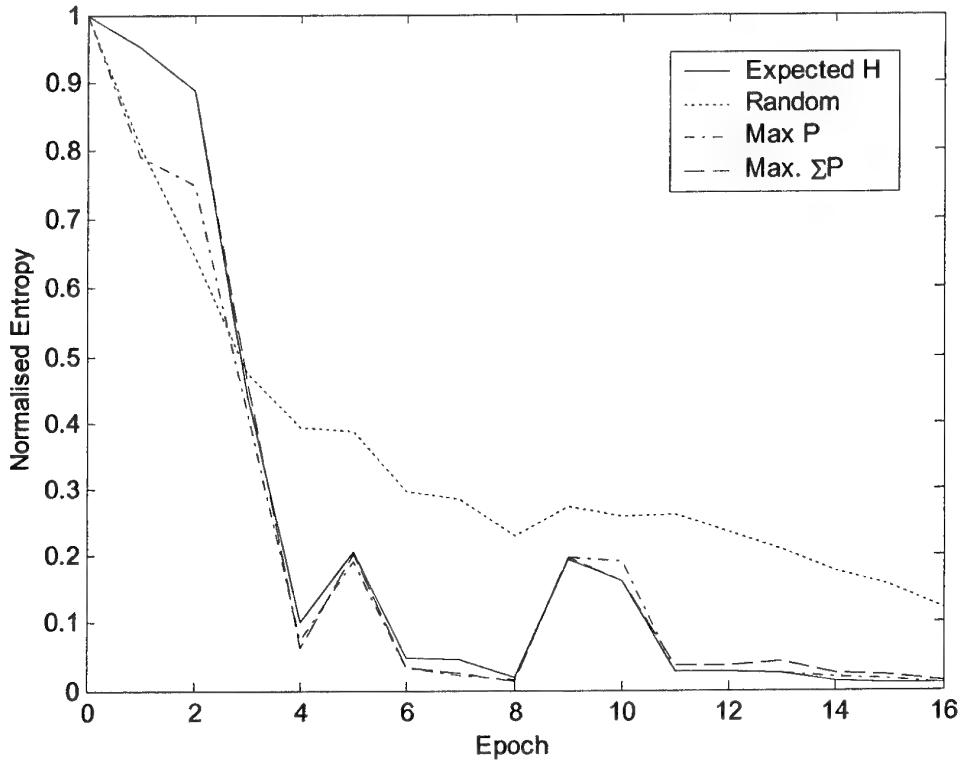


Figure 5: Comparison of the swath selection techniques for the target motion path shown in Figure 2, for $\xi = 12.5$ and $p_d = 0.95$, $p_{fa} = 0.001$.

Figure 6 shows a comparison between the four techniques for the motion shown in Figure 3 with parameter values of $\xi = 30$, $p_d = 0.95$ and $p_{fa} = 0.001$. Once again the random or baseline technique is overall the worst, whereas the other three techniques are generally the same and reach almost zero entropy after 6 epochs. The random and maximum probability selection techniques are better for small epochs, for the reasons discussed above. There is also a spike in the entropy at epoch 5 when the target distribution is distributed over swaths 1 and 2. The three techniques are able to lock onto the target and reduce the entropy to almost zero. There are also small spikes at

epoch 10 when the target changes direction, and epoch 13 when the probability is spread over two swaths. Nevertheless, the increase in entropy is very small.

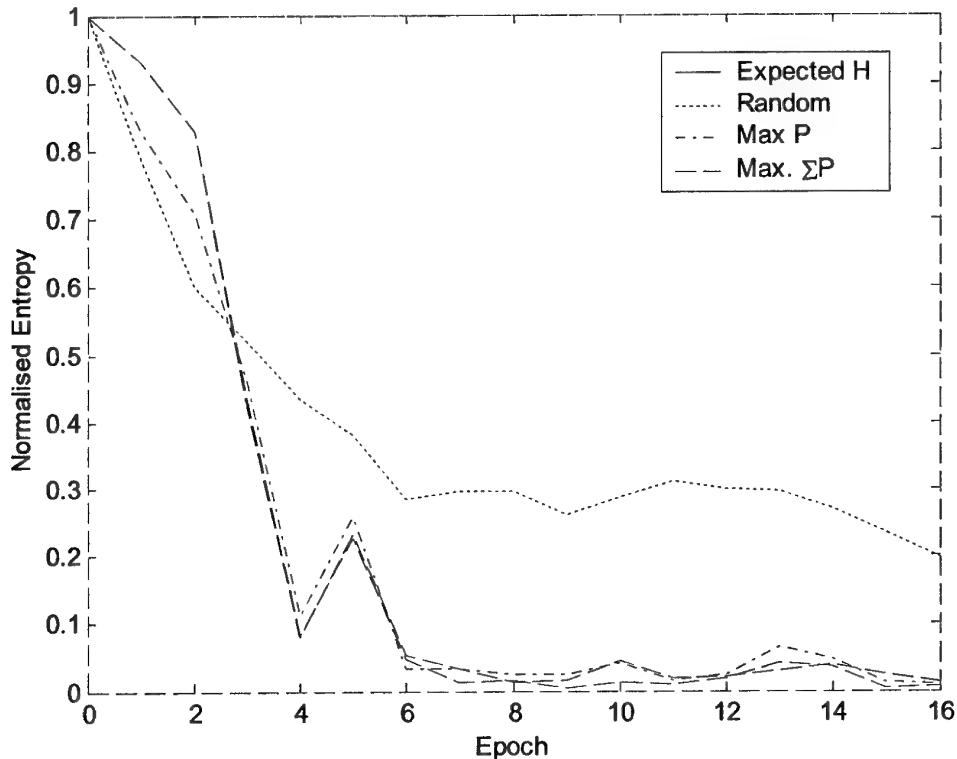


Figure 6: Comparison of the swath selection techniques for the target motion path shown in Figure 3 for $\xi = 30$ and $p_d = 0.95$, $p_{fa} = 0.001$.

Finally, Figure 7 is for a similar scenario to that of Figure 6, but with a more non-ideal sensor: $p_d = 0.80$ and $p_{fa} = 0.01$. Again the random swath selection, or baseline, technique is the worst, while the other three are very similar. In this case there is a broad spike visible at epoch 13 when the target's probability is distributed over swaths 2 and 3.

The very non-ideal nature of the sensor and the broad Gaussian motion model ($\xi = 30$), results in the techniques taking many epochs to reduce the entropy and even then the minimum is not close to zero.

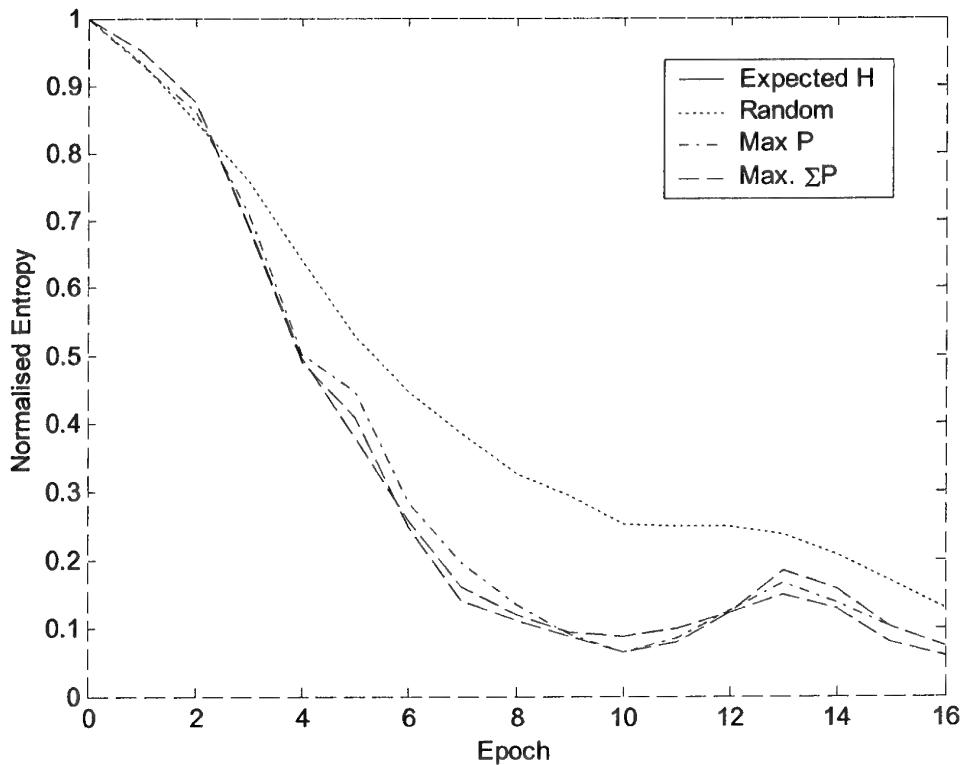


Figure 7: Comparison of the swath selection techniques for the target motion path shown in Figure 3, for $\xi = 30$ and $p_d = 0.80$, $p_{fa} = 0.01$.

The results above indicate to us a number of points: firstly, the expected entropy algorithm developed in previous sections of this report can track a target, albeit in a simplified scenario for the time being. Secondly, selecting the optimal swath based on the expected entropy gives results as good as a number of other intuitive techniques, and much better than the baseline performance of randomly selecting a swath. Although more complicated than the other two, it is the only technique that can take into account the non-ideal nature of the sensors.

11. Demonstration of Benefits for Analytically Tractable Example

The inconclusiveness of the previous results has prompted some analysis of an analytically tractable example. The fundamental question to be answered is whether, for a given *a priori* probability distribution, the use of expected entropy as a measure of surveillance performance for optimising resource allocation for space-based

surveillance assets causes an extremal choice of sensor resolution (such as maximum area / minimum resolution or minimum area / maximum resolution) that could be arrived at by the application of simple rules, or an intermediate choice, which requires computation to determine.

Fundamentally, the balance between the second and third terms in the r.h.s represents the trade-off between sensor area coverage and sensor resolution of the expected entropy expression:

$$h' = h - \sum_i p_i \sum_l L(l|i) \log L(l|i) + \sum_l P(l) \log P(l)$$

Note that the second term is positive and has a tendency to increase entropy whereas the third term is negative and has a tendency to reduce it. The simplest way of exploring this issue is in the localised formulation because each cell has only two states (making the likelihood function and entropy computation easy) and the entropy for the joint distribution over all cells is simply the sum of the entropies for the individual cells. This can be explored without using simulation by assuming a prior distribution and computing the expected entropy for a choice of sensor resolutions (and corresponding coverages). Note that in the localised formulation, cells are independent so the result of inspecting cells does not affect cells not inspected. In other words, the total change in expected entropy only involves a summation over those cells that are inspected.

In the remainder of this section we illustrate the potential of the developed algorithm to select an "optimal" swath mode, or position, from a suite of on-board sensors. One would hope that the swath chosen by the algorithm, which is the swath that minimises the entropy over the ROI, results in a natural trade-off between resolution (cell area) and overall swath size. Ideally, the resolution should be high as possible i.e., small cell area, thereby increasing the probability of detection and decreasing the false alarm rate (FAR). Furthermore, a large swath size over the ROI is desirable, as this increases the likelihood that the target lays within the swath during a satellite overpass.

Generally these two desirable features cannot be achieved simultaneously with most sensors i.e., selecting a sensor with high resolution, or small cell area, results in a smaller swath area than a sensor with coarser resolution. From the specifications in "RADARSAT Illuminated" this appears to be a reasonable assumption. Thus, there is a trade-off between the two, which we expect the algorithm to determine in some way.

To illustrate this consider the following simple example of choosing between four of RADARSAT's modes: H1, S1, W1 and L1. The Table below gives their cell area, swath width and area, and the approximate total number of cells within the swath.

Table 2: Details of the four sensor modes used (see "RADARSAT Illustrated").

Mode	Cell Area (m ²)	Swath Width (km)	Swath Area (m ²)	Total Cells
H1	535	75	5.63×10^9	1.052×10^7
S1	702	100	1.00×10^{10}	1.425×10^7
W1	959	165	2.72×10^{10}	2.840×10^7
L1	980	170	2.89×10^{10}	2.949×10^7

The following two figures show the swath width (SW) and total number cells (N) plotted against the cell area (A) as circular points. Two quadratic curves are fitted to the data, resulting in the following relationship:

$$SW = 0.0002315A^2 - 0.1354431A + 81.1669018 \text{ km,}$$

$$N = 73.619A^2 - 68344.966A + 25998324.264.$$

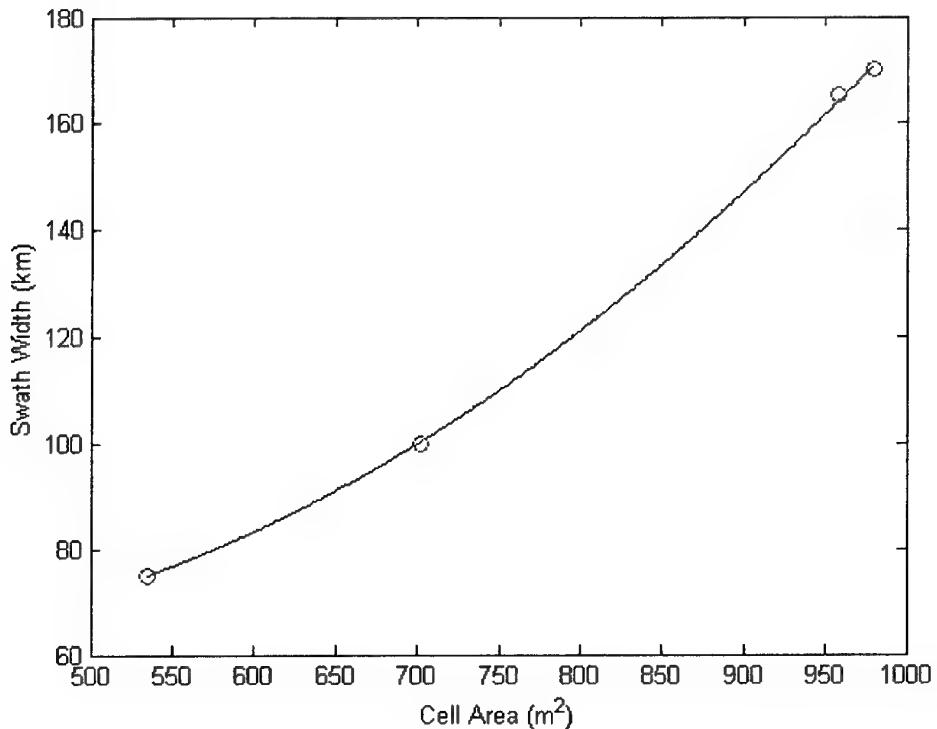


Figure 8: Data and fitted curve of 'Swath Width' versus 'Cell Area'.

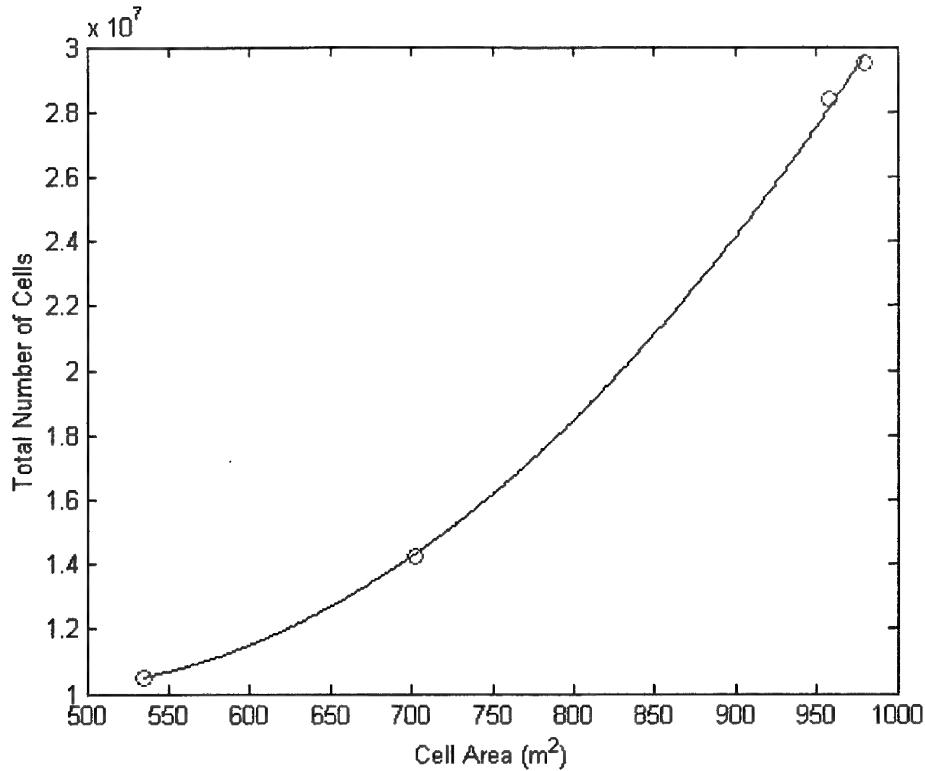


Figure 9: Data and fitted curve of 'Total Number of Cells' versus 'Cell Area'.

To determine how each sensor mode or swath reduces the total entropy, one needs to compute the change in entropy per cell, once the probabilities of detection and false alarm are estimated. We take a simple approach, which allows us to illustrate the principle without delving unnecessarily into the complex modeling aspects of sensors and the radar processing of the returns.

We assume that the probability of detection is a function of the swath cell area covered by the target. The probability of detection is computed from a sigmoid function with the free variable being the fraction of the cell area covered by the target:

$$u = \left(\frac{2A_T}{A} \right) - 1,$$

$$p_d(u) = \frac{1}{1 + e^{-\alpha u}},$$

where A_T is the target's area, A is the swath's cell area, and α is the slope of the function at $u = 0$.

For this experiment a value of $\alpha = 5$ was used. The FAR is based upon the cell area multiplied by the density function of the probability of false alarm. If a cell has area A_0 and FAR of p_{fa0} , then the FAR of a cell with area A is given by

$$p_{fa} = \frac{A}{A_0} p_{fa0}.$$

In our case, a nominal FAR of 10^{-8} was assumed for a cell area of 500 m^2 , and so the FAR for the four modes were computed from their relative cell area values. The probabilities of detection and false alarm for the four sensor modes are given below for a target area of 600 m^2 .

Table 3: Computed values of probabilities of detection and false alarm.

Mode	p_d	p_{fa}
H1	.9980	1.068×10^{-8}
S1	.9720	1.404×10^{-8}
W1	.7796	1.916×10^{-8}
L1	.7544	1.960×10^{-8}

Once p_d and p_{fa} are estimated, the change in entropy per cell can be computed using equation 6.1 from Section 6.4, assuming that the probability of a target existing in a cell is $p = 0.70$. The total entropy change in any given swath is the total entropy change per cell multiplied by the number of cells in that swath, and its magnitude is denoted by $N|\Delta h|$.

The figure below shows a 3-D line plot of how the swath width and total entropy vary as the cell area varies. The four circles represent the actual data. The solid line is the total entropy change, while the dotted line is the projection of this onto the 'Swath Width' and 'Area' plane. The solid line is the plot of the values obtained from the fitted curves described above.

Although there are only four points, for the particular parameters and approximations chosen, there is a peak in the total entropy change curve. The entropy method tells us that the "optimal" cell area is somewhere between 850 m^2 and 900 m^2 . Of course, the four RADARSAT modes do not have a cell area exactly equal to this, so the mode with the closest cell area is chosen, in this case mode W1.

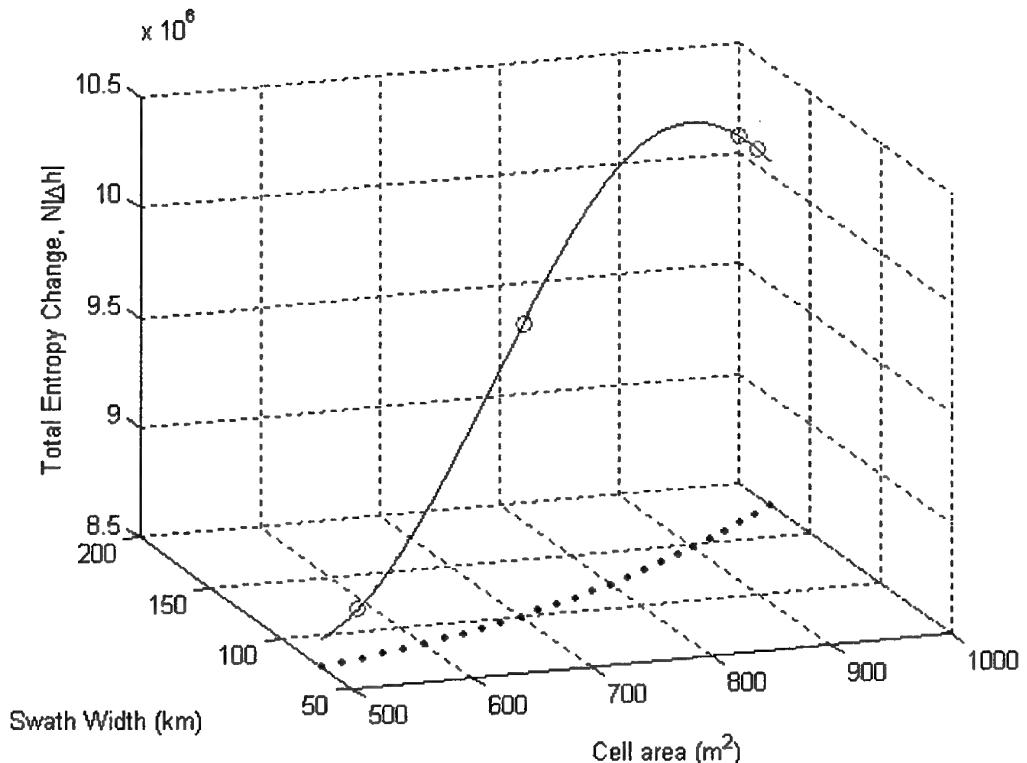


Figure 10: Plot of 'Total Entropy Change' vs. 'Swath Width' and 'Cell Area'. The dotted line a projection of the solid line onto the 'Swath Width' and 'Cell Area' place.

12. Conclusions

This report provides optimal solutions to the problem of dynamically determining the allocation and control of space-based surveillance resources for the purpose of detecting, locating and tracking surface maritime targets using the GAMBIT formalism [Berry & Fogg (1)]. The sensors' performances were modelled in terms of their probabilities of detection and false alarm. Three types of target motion model were considered, deterministic, Gauss-Markov and conservative. Multiple targets were allowed. The solution progressed by cyclical target motion prediction and Bayesian update. Significant computational efficiency was achieved by use of a 'localised formulation' of the problem.

The approach presented provides a rational basis for space-based surveillance asset optimisation which sets information as the objective, incorporates platform dynamics and constraints, sensor capabilities, stochastic effects and control variables within a single formulation.

As well as being an issue of interest in its own right, this is an instance of the more general surveillance asset allocation problem.

To illustrate the implementation of this process, the use of SAR satellites with several modes allowing different swaths and resolutions was considered. Optimal selection of swaths for target tracking, by choosing the swath for each access that minimizes the expected entropy (MOE) across the region of interest, was demonstrated.

Three other MOE were used to select the swath, random, maximum probability, and maximum sum of probability. All were shown to be better than fixed swath or random swath selection.

It might be expected that with less ideal targets which change speed or direction more frequently or dramatically, a Monte Carlo simulation involving multiple runs would be more discerning statistically.

It was shown that situations arise in which an intermediate rather than extreme choice of sensor resolution do result in an optimal solution of the surveillance asset resource allocation problem, thereby providing justification for the proposed approach.

13. Recommendations

Further work is required to fully validate the concepts reported here, both in terms of confirming and quantifying the benefits to space-based surveillance effectiveness and in demonstrating the practicality of implementation. In order to achieve this, a detailed simulation model is required which captures satellite dynamics, ideally based upon a high fidelity model (such as Satellite Tool Kit), and a high fidelity SAR sensor model (such as SBST [9]) is also needed. For more general and complex surveillance asset allocation problems, which potentially involve constrained optimisation of a mix of discrete and continuous parameters over time, commercial optimisation software (such as ILOG) should be investigated.

14. References

- [1] Berry, P.E. and Fogg, D.A.B., "GAMBIT: Gauss-Markov and Bayesian Inference Technique for Information Uncertainty and Decision-making in Surveillance Simulations", DSTO Research Report in draft 2003.
- [2] Stone L.D., Barlow C.A. and Corwin T.L., "Bayesian Multiple Target Tracking", Artech House, 1999.
- [3] Kastella K., "A Microdensity Approach to Multitarget Tracking", Proceedings of the 3rd International Conference on Information Fusion, France, July 10-13 2000.
- [4] Xianxing L., Shengwei L., Pan Q. and Zhang H., "A Method of Sensor Management Based on a Probability Statistical Model", (in Chinese) Control Theory and Applications, Vol 18, No 5, October 2001.
- [5] Berry P.E., "A Dynamic Asset Tasking Technique for Integrated Surveillance Operations", DSTO-RR-0246, 2002.
- [6] Fogg D., "Entropy-based Measures of Effectiveness for Integrated Surveillance", Defence Operations Research Conference, 17-18 March 1999.
- [7] Xianxing L., Shen S., Pan Q. and Zhang H., "An Algorithm of Sensor Management Based on Information Entropy", (in Chinese) Acta Electronica Sinica, Vol 28, No 9, September 2000.
- [8] RADARSAT International, "RADARSAT Illuminated", Preliminary Version, July 1995.
- [9] Van Antwerpen C., Kruizing E., "Space Based Synthetic Aperture Radar Tool-A Code and Model for Estimating the Maritime Detection Capability of Generic Space-Borne Synthetic Aperture Radars", DSTO-TR-1050.
- [10]: Salmon, D.J., "The Characteristics of the Second-Order Target Model", Royal Aircraft Establishment Technical Report 85071, August 1985.

15. Mathematical Notation

$A, (A_n)$	Set of cells within \mathfrak{R} that are accessed by a swath (at epoch n).
ΔA	Area covered by each cell in \mathfrak{R} .
$c_I(\cdot)$	Likelihood $L_I(\cdot)$ summed over all possible detection events.
$c_H(\cdot)$	Likelihood $L_H(\cdot)$ summed over all possible detection events.
$c_{III}(\cdot)$	Likelihood $L_{III}(\cdot)$ summed over all possible detection events.
${}^M C_k$	Combinations of taking k from M .
$E(\cdot)$	Statistical expectation value.
h'	Expected entropy.
h	Measured entropy.
h_i^n	Entropy of cell i at the n^{th} epoch.
$h_i^n(0)$	Entropy of cell i at the n^{th} epoch after a non-detection event.
$h_i^n(1)$	Entropy of cell i at the n^{th} epoch after a detection event.
\bar{h}_i^n	Entropy of cell i at the n^{th} epoch (immediately before $(n+1)^{\text{th}}$ epoch).
k	Number of detection events.
\underline{l}	Set of observations (detections) $\{j_1, j_2, \dots, j_k\}$ at epoch n .
$L(\cdot)$	Likelihood function.
$L(\underline{l}/i)$	Likelihood of obtaining measurement set \underline{l} given the target is in cell i .
$L_I(\cdot)$	Likelihood function for a cell accessed by swath but with no detection event.
$L_H(\cdot)$	Likelihood function for a cell accessed by swath and with a detection event.
$L_{III}(\cdot)$	Likelihood function for a cell not accessed by swath.
M	Number of cells accessed by swath A .
m	Number of fixed time intervals Δt between successive accesses.
m_n	Number of fixed time intervals Δt between epochs n and $(n-1)$.
N	Number of cells in the region of interest, \mathfrak{R} .
n	Epoch number.
<i>p.d.f.</i>	Probability density function.
p_A	Sum of probabilities of all cells under the swath.
p_d	Probability of detection.
p_{fa}	Probability of false alarm.
$p_i(t)$	Probability of a target being in cell i at time t .
p_i^n	<i>A posteriori</i> probability distribution of target in cell i at epoch n .
\bar{p}_i^n	<i>A priori</i> probability distribution of target in cell i at epoch n (time t_n).

$p_{i,j}^n$	-	Probability that target moved from cell i at epoch $(n-1)$ to cell j at epoch n .
$\bar{p}_{i,j}^n$	-	<i>A priori</i> probability distribution of target moving from cell i at epoch $(n-1)$ to cell j at epoch n .
P_{fa}	-	False alarm rate per unit area.
$P(\underline{l})$	-	Probability of obtaining the measurement set \underline{l} .
$\text{Pr}\{\cdot\}$	-	Generic probability distribution.
$P_z(z)$	-	p.d.f. of random variable z .
$P_\theta(\theta)$	-	p.d.f. of random variable θ .
$Q(i, j, t_{n+1} - t_n)$	-	Probability that target is in cell j at time t_{n+1} given that it was in cell i at time t_n .
\mathfrak{R}	-	Region of interest.
s	-	Maximum target velocity.
t	-	Time from epoch 0.
t_n	-	Time of n^{th} satellite access to the region of interest.
Δt	-	Fixed time interval.
T	-	Time between epochs t_n and t_{n+1} .
x, y	-	Cartesian coordinates.
x_m, y_m	-	Coordinates after m time steps Δt . Equivalent to $x(m\Delta t), y(m\Delta t)$.
z	-	RV for the distance travelled by target before course change.
α	-	Sensor look angle.
λ	-	Average rate of course changes per unit distance.
ε_m	-	Sequence of target's acceleration.
ξ	-	Variance of target's acceleration.
μ_x, μ_y	-	Target's expected position in Cartesian coordinates.
σ	-	Standard deviation.
Λ	-	Set of all possible observations, \underline{l} .
θ	-	Random variable for the target's course change.
$\{j_1, j_2, \dots, j_k\}$	-	Set of k detection events.

Appendix A: Derivation of a general expression for the expected entropy over a set of observations

Assuming that there is a single target in \mathfrak{R} and that observation in a swath A results in a set of measurements $\underline{l} = \{j_1, j_2, \dots, j_k\}$ then Bayes' rule is used to update the probability distribution for the target location thus

$$p_i' = \frac{L(\underline{l} | i) p_i}{P(\underline{l})}$$

where $P(\underline{l}) = \sum_{i \in \mathfrak{R}} L(\underline{l} | i) p_i$ is the probability of the set of measurements $\underline{l} = \{j_1, j_2, \dots, j_k\}$ occurring. Let the set of all possible measurements \underline{l} be Λ then the entropy before a set of measurements $\underline{l} \in \Lambda$ is

$$h = -\sum_{i \in \mathfrak{R}} p_i \log p_i$$

and afterwards is

$$\begin{aligned} h'(\underline{l}) &= -\sum_{i \in \mathfrak{R}} p_i' \log p_i' \\ &= -\sum_i \left(\frac{L(\underline{l} | i) p_i}{P(\underline{l})} \right) \log \left(\frac{L(\underline{l} | i) p_i}{P(\underline{l})} \right) \\ &= -\frac{1}{P(\underline{l})} \left[\sum_{i \in \mathfrak{R}} p_i L(\underline{l} | i) \log L(\underline{l} | i) + \sum_{i \in \mathfrak{R}} L(\underline{l} | i) p_i \log p_i \right] + \log P(\underline{l}) \end{aligned}$$

Hence the expected entropy over all possible sets of measurements $\underline{l} \in \Lambda$ is

$$\begin{aligned} h' &= E_{\underline{l} \in \Lambda} \{h'(\underline{l})\} \\ &= \sum_{\underline{l} \in \Lambda} P(\underline{l}) h'(\underline{l}) \\ &= h - \sum_{i \in \mathfrak{R}} p_i \sum_{\underline{l} \in \Lambda} L(\underline{l} | i) \log L(\underline{l} | i) + \sum_{\underline{l} \in \Lambda} P(\underline{l}) \log P(\underline{l}) \end{aligned}$$

after simplification.

Appendix B: Interpretation of the expression for the expected entropy

With

$$h' = h - \sum_i p_i \sum_l L(l|i) \log L(l|i) + \sum_l P(l) \log P(l)$$

the second term on the right hand side represents the *increase* in entropy arising from the imperfections in the sensors. This is because if the likelihood functions are all 1's or 0's then the term vanishes. This would correspond to the case of perfect sensors. So this term and its sign are explainable.

As for the last term on the r.h.s., this obviously represents a *decrease* in entropy because of its sign. We can understand this by considering the case of a perfect sensor that can see everywhere (i.e. $p_d = 1$ and $p_{fa} = 0$). Then the observation label l corresponds to a cell i and we see that $\sum_l P(l) \log P(l) = \sum_i p_i \log p_i = -h$ and then $h' = 0$. This says that

the sensor will 'see' a target in a cell in proportion to the probability of it being there and the entropy immediately drops to zero (perfect information). Thus, the sign seems to make sense.

Now consider the case of a perfect sensor, which cannot see everywhere, but just sees a subset of all the cells (a swath). We want to choose to look in those cells that make the biggest contribution to the reduction in h . We can only sum over the states in the observation space in $\sum_l P(l) \log P(l)$ i.e., over the cells in which we look. This is now

closer to the approach of summing probabilities over the cells and choosing the swath with the biggest sum, or choosing the swath that contains the cell with the biggest probability. Thus, we choose the swath S that makes $\sum_{i \in S} p_i \log p_i$ the most negative and hence reduces the expected entropy the most.

An interesting question is whether maximising the change in entropy by choosing the cells that provide the biggest $p \log p$ contributions is equivalent to maximising p or $\sum p$. Note that $p \log p$ is biggest when p is closest to 0.5, however there is a constraint that all the p 's must sum to 1 so it is not obvious in general.

Unlike the approach of maximising probabilities to find the optimum swath, this entropy approach automatically takes account of the imperfections in the sensors because of the second term on the r.h.s. (eg multipath effects, cloud cover for EO sensors) and hence is more general.

Appendix C: Implementation Procedures

C.1. Discretized Tracking Procedure

The following procedure explains the discretized tracking implementation, which is presented in Section 9.

1. Initialise search and tracking:
 - Epoch: $n = 0$
 - Cell probabilities: $p_{m,i}^n = \frac{1}{N}, i = 1, 2, \dots, N$
2. Compute target motion transition probabilities:
 - If $n = 0$: compute $Q(j/i, m)$, $m = i$ using algorithm in Appendix C.2.
 - Else: compute $Q(j/i, m)$, using algorithm in Appendix C.2.
3. Target motion update:
 - $\bar{p}_{i,j}^{n+1} = \sum_m Q(j/i, m) p_{m,i}^n$
 - $\bar{p}_j^{n+1} = \sum_i \bar{p}_{i,j}^{n+1}$
4. Select optimal swath by minimising the expected entropy. See Appendix C.3.
5. Update target's position.
6. Generate detection events.
7. Bayesian update. See Appendix C.5.
 - $p_{i,j}^{n+1} = B\{\bar{p}_{i,j}^{n+1}\}$
8. Compute overall target probability:
 - $p_j^{n+1} = \sum_i p_{i,j}^{n+1}$
9. Last epoch of simulation?
 - If *yes*, then END.
 - If *no*, then $n = n + 1$ and go to 2.

C.2. Gauss-Markov Motion Update Procedure

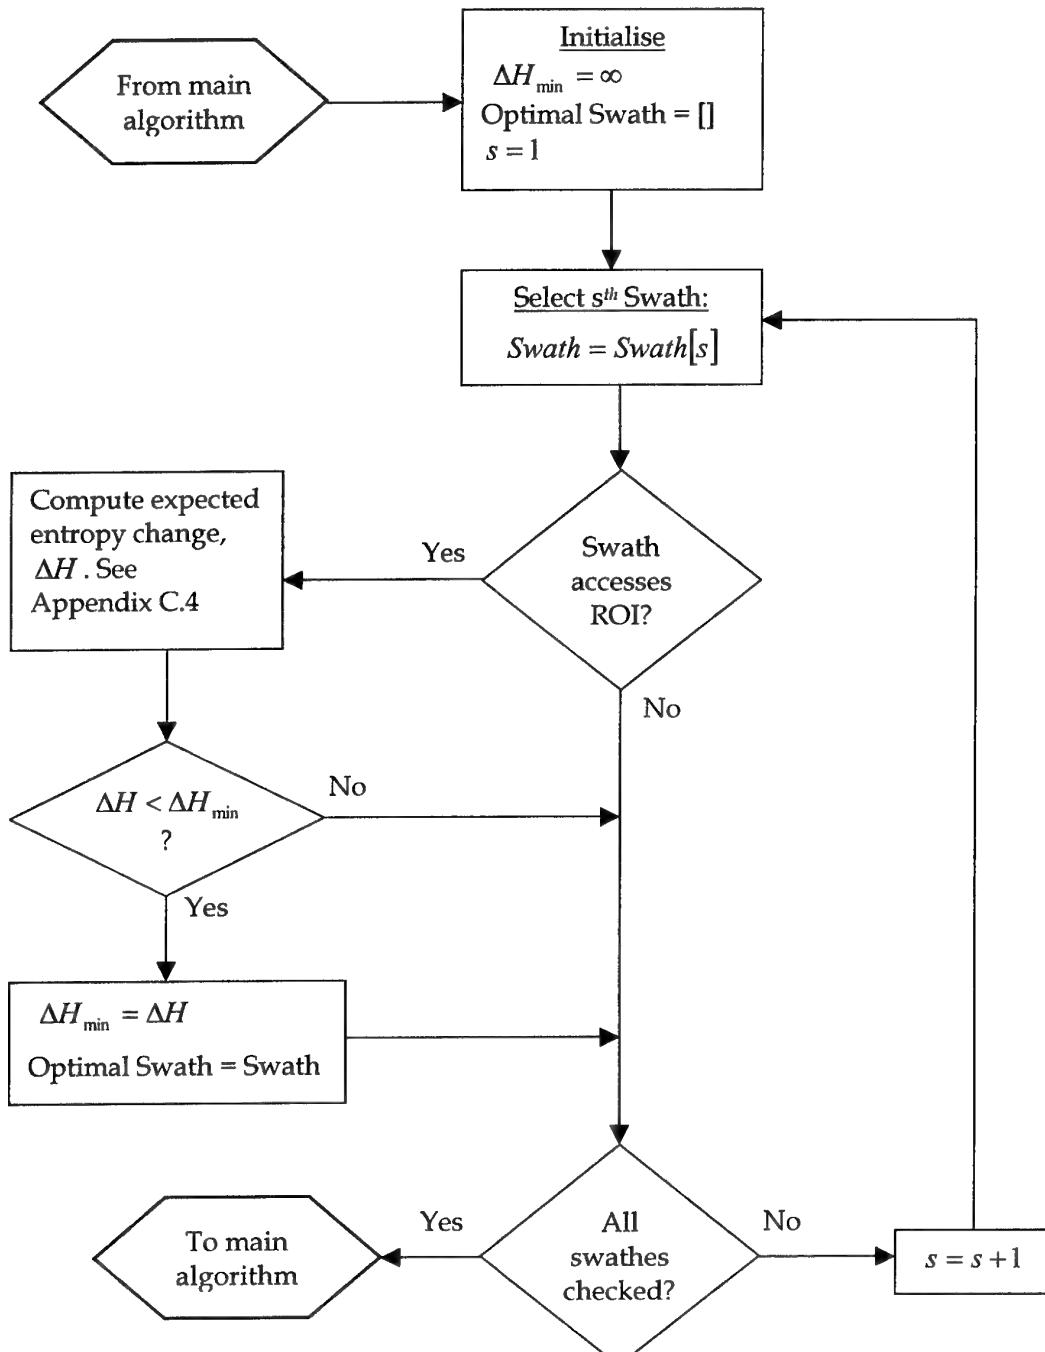
The following procedure details the implementation of the Gauss-Markov motion model as described in Section 9.1.2. It is called by the main tracking procedure in Appendix C.1, and then returns to it once it has completed computing the probabilities.

1. Estimate Velocity.
 - Estimate target's velocity $v_{m,i}^n$ i.e., target moved from cell m at epoch $(n-1)$ to cell i at epoch n .
2. Predict Target's Next Position.
 - Predict target's position at epoch $(n+1)$ assuming velocity $v = v_{m,i}^n$. Let the cell's x and y coordinates be $(\hat{x}^{n+1}, \hat{y}^{n+1})$.
3. Compute Variance.
 - Compute the variance, σ^2 , of the target's acceleration.
4. Compute Probabilities.
 - Let the j^{th} cell's x and y indices be (x_j, y_j) . Compute the integral:

$$Q(j/i, m) = \frac{1}{2\pi\sigma^2} \int_{y_j-0.5}^{y_j+0.5} \int_{x_j-0.5}^{x_j+0.5} \exp\left[-\frac{1}{2\sigma^2} \left\{ (x - \hat{x}^{n+1})^2 + (y - \hat{y}^{n+1})^2 \right\}\right] dx dy$$
5. Return to Main Algorithm. (i.e., Return to 2. in main algorithm in Appendix C.1).

C.3. Optimal Swath Selection Procedure

The following procedure presents the logic in computing the optimal swath that reduces the expected entropy; see Section 6.3. This procedure is called by the tracking procedure of Appendix C.1, and returns to it once completed. It calls the procedure to compute the expected entropy change, as described in Appendix C.4.



C.4. Expected Entropy Change Procedure

The following procedure details the process of computing the expected entropy change as described in Section 6.4. It is called by the optimal swath selection procedure in Appendix C.3, and also returns to it once the expected entropy change is calculated.

1. Assign Variables.

- M - no. of cells in region (A) covered by swath
- p_d - Swath's probability of detection
- p_{fa} - Swath's false alarm rate
- $p_A = \sum_{c \in A} p_c$

2. Compute Probabilities

- $c_I = p_A \sum_{k=0}^{M-1} {}^{M-1}C_k (1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1} \log \{ (1-p_d) p_{fa}^k (1-p_{fa})^{M-k-1} \}$
- $c_{II} = p_A \sum_{k=1}^M {}^M C_{k-1} p_d p_{fa}^{k-1} (1-p_{fa})^{M-k} \log \{ p_d p_{fa}^{k-1} (1-p_{fa})^{M-k} \}$
- $c_{III} = (1-p_A) \sum_{k=0}^M {}^M C_k p_{fa}^k (1-p_{fa})^{M-k} \log \{ p_{fa}^k (1-p_{fa})^{M-k} \}$
- $P_0 = (1-p_d)(1-p_{fa})^{M-1} p_A + (1-p_{fa})^M (1-p_A)$
- $P_M = p_d p_{fa}^{M-1} p_A + p_{fa}^M (1-p_A)$
- $P_k = p_{fa}^{k-1} (1-p_{fa})^{M-k-1} \left[(1-p_d) p_{fa} \left(p_A - \sum_{l=1}^k p_{j_l} \right) + p_d p_{fa} (1-p_{fa}) \sum_{l=1}^k p_{j_l} + (1-p_A) p_{fa} (1-p_{fa}) \right]$
- $P = P_0 \log P_0 + P_M \log P_M + \sum_{k=1}^{M-1} \sum_{\{j_1, j_2, \dots, j_k\}} P_k \log P_k$

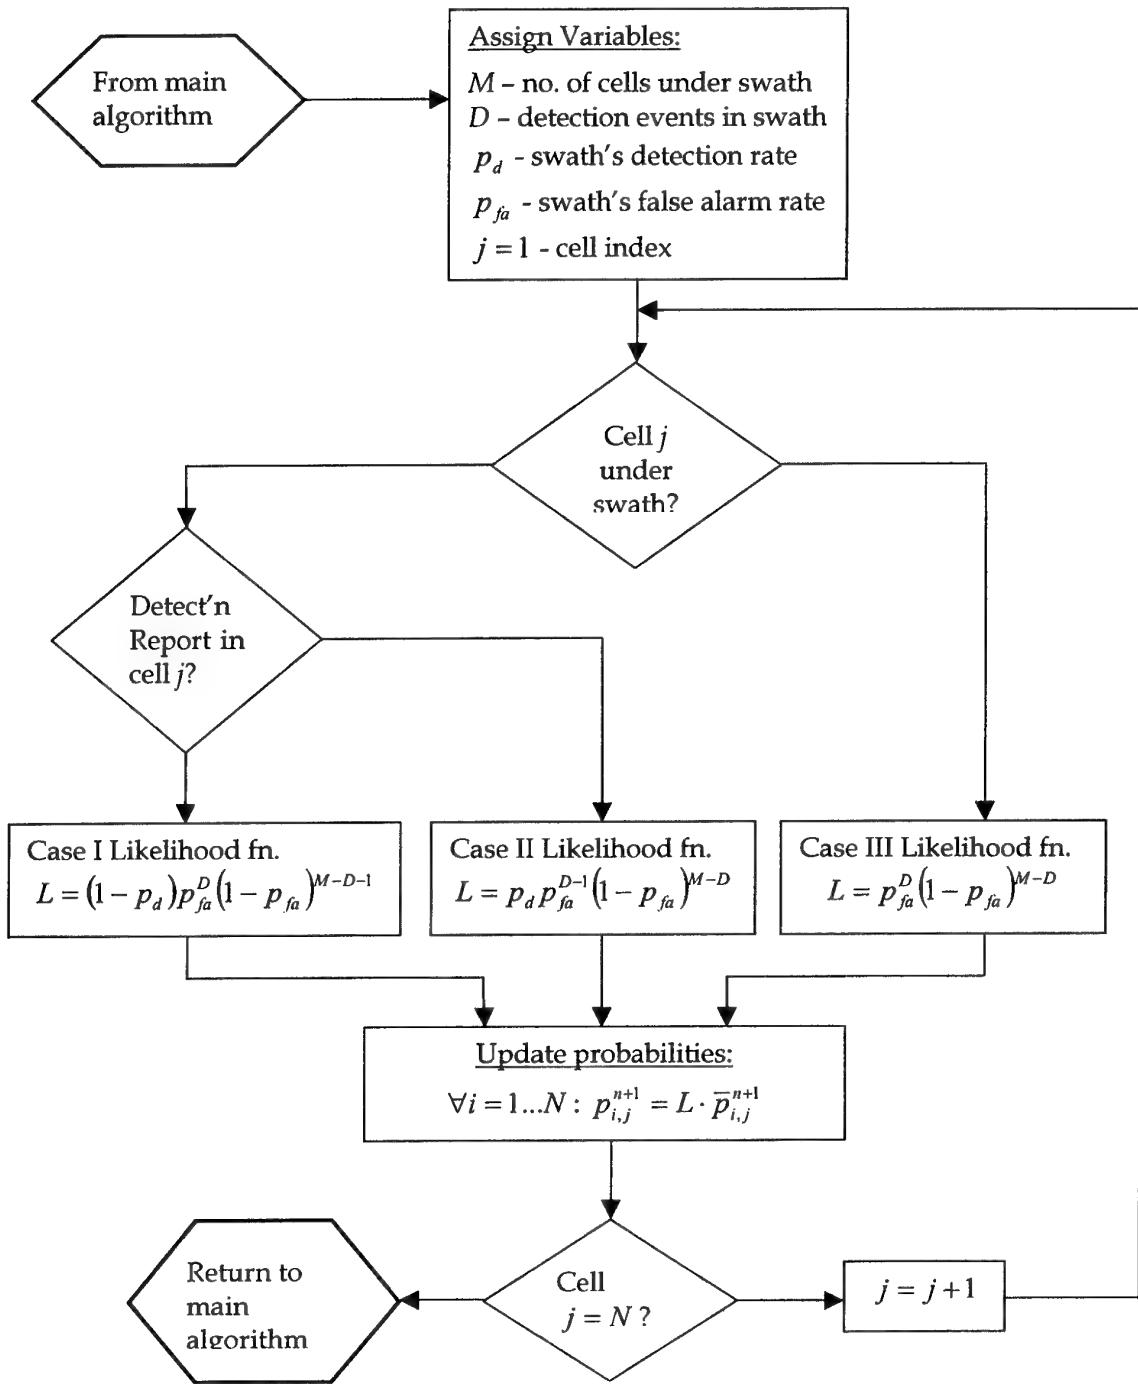
3. Compute Entropy Change

- $\Delta H = -c_I - c_{II} - c_{III} + P$

4. Return to Optimal Swath Selection Procedure in Appendix C.3.

C.5. Bayesian Update Procedure

This procedure explains the Bayesian update process as described in Sections 6.4 and 7.1. It is called by the main tracking algorithm of Appendix C.1, and also returns to it once completed.



DISTRIBUTION LIST

Optimal Search, Location and Tracking of Surface Maritime Targets by a
Constellation of Surveillance Satellites

P. Berry, C. Pontecorvo and D.A.B. Fogg

AUSTRALIA

DEFENCE ORGANISATION

Task Sponsor

DGC4ISREW

S&T Program

Chief Defence Scientist

FAS Science Policy

AS Science Corporate Management

Director General Science Policy Development

Counsellor Defence Science, London (Doc Data Sheet)

Counsellor Defence Science, Washington (Doc Data Sheet)

Scientific Adviser to MRDC Thailand (Doc Data Sheet)

Scientific Adviser Joint

Navy Scientific Adviser (Doc Data Sheet and distribution list only)

Scientific Adviser - Army (Doc Data Sheet and distribution list only)

Air Force Scientific Adviser

Director Trials

} shared copy

Information Sciences Laboratory

Chief of ISR Division (Doc Data Sheet and Distribution List Only)

Research Leader WAS (Doc Data Sheet and Distribution List Only)

Head SSMA

P. Berry

C. Pontecorvo

DSTO Library and Archives

Library Edinburgh 2 copies

Australian Archives

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet only)

Director General Land Development

Director General Aerospace Development (Doc Data Sheet only)

Knowledge Staff

Director General Command, Control, Communications and Computers (DGC4)

(Doc Data Sheet only)

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW (Doc Data Sheet and distribution list only)

Army

ABCA National Standardisation Officer, Land Warfare Development Sector, Puckapunyal (4 copies)

SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD (Doc Data Sheet only)

NPOC QWG Engineer NBCD Combat Development Wing, Puckapunyal, VIC (Doc Data Sheet relating to NBCD matters only)

Intelligence Program

DGSTA Defence Intelligence Organisation Manager, Information Centre, Defence Intelligence Organisation

Defence Libraries

Library Manager, DLS-Canberra
Library Manager, DLS - Sydney West (Doc Data Sheet Only)

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong, VIC
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

OTHER ORGANISATIONS

National Library of Australia
NASA (Canberra)
State Library of South Australia

OUTSIDE AUSTRALIA

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US
Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK

SPARES (10 copies)

Total number of copies: 47

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA		1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)		
2. TITLE Optimal Search, Location and Tracking of Surface Maritime Targets by a Constellation of Surveillance Satellites (U)		3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Paul E. Berry, Carmine Pontecorvo, David A.B. Fogg		5. CORPORATE AUTHOR Information Sciences Laboratory PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1480	6b. AR NUMBER AR-012-869	6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE <i>Aug 2003</i>	
8. FILE NUMBER 9505-23-98	9. TASK NUMBER JTW 00/255	10. TASK SPONSOR DGC4ISREW	11. NO. OF PAGES 54	12. NO. OF REFERENCES 10
13. URL on the World Wide Web			14. RELEASE AUTHORITY	
http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1480.pdf			Chief, Intelligence, Surveillance and Reconnaissance Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>				
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111				
16. DELIBERATE ANNOUNCEMENT				
No Limitations				
17. CITATION IN OTHER DOCUMENTS		Yes		
18. DEFTEST DESCRIPTORS				
Space surveillance (spaceborne), integrated systems, ocean surveillance, target acquisition, target tracking.				
19. ABSTRACT				
The issues associated with the maximal exploitation of space-based surveillance resources are unique due to the nature of the platforms, their sensors and intermittent communications links to the ground station for target information processing and sensor tasking. This report applies the GAMBIT formalism [Berry & Fogg (1)] to the networked sensor decision problem of determining the optimal allocation of space-based surveillance resources for the purpose of detecting, locating and tracking surface maritime targets. As well as being an issue of interest in its own right, this application is an instance of the more general surveillance asset allocation problem.				